NGAIZA SERIES OF ADVANCED PHYSICS

We let the physics dictate the book. We never let the marketing dictate the physics.

ENJOY & MASTER ADVANCED PHYSICS

PHYSICS INTERACTIONS

Heat, Waves & Electricity

Ngaiza Books App
Scan to download on Google Play

Scan for interactive solutions, extra questions, and all Ngaiza titles on your phone.

Enjoy & Master Advanced Physics
Physics Interactions: Heat, Wave & Electricity

Ngaiza Series of Advanced Physics

© 2026 Ngaiza Education Hub. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means: electronic, mechanical, photocopying, recording, or otherwise without the prior written permission of the publisher, except for brief quotations in reviews or academic commentary.

First Edition, 2026

ISBN 978-9987-9677-5-9

Writers
Mr. Amir Juma, Mr. Kelvin Mashu, Mr. Endrew Mariki, and Mr. Morris Mulokozi

Editors
Dr. Aloyce Julian and Dr. Alex Muganyizi

Series Creator and Director
Ngaiza Lusima

Published by
Ngaiza Education Hub
P.O. Box 60791
Mbezi Beach, Dar es Salaam, Tanzania

www.ngaizabooks.com
contact@ngaizabooks.com

Every effort has been made to ensure the accuracy of the content in this book. The publisher welcomes corrections and suggestions for future editions.

TABLE OF CONTENTS

Preface i

The Map of Interactions iii

PART ONE: PROPERTIES OF MATTER AND THERMAL PHYSICS

Chapter 1 Elasticity 2

Miscellaneous Worked Examples 48

Digging Deeper Exercise 1 57

Chapter 2 Temperature and Heat 60

Miscellaneous Worked Examples 97

Digging Deeper Exercise 2 103

Chapter 3 Heat Transfer 107

Miscellaneous Worked Examples 158

Digging Deeper Exercise 3 164

Chapter 4 Kinetic Theory of Gases 169

Miscellaneous Worked Examples 197

Digging Deeper Exercise 4 205

Chapter 5 Thermodynamics 208

Miscellaneous Worked Examples 261

Digging Deeper Exercise 5 270

Examination Questions for Part One 274

PART TWO: VIBRATIONS AND WAVES

Chapter 6 Simple Harmonic Motion 284

Miscellaneous Worked Examples 350

Digging Deeper Exercise 6 357

Chapter 7 Wave Motion 360

Miscellaneous Worked Examples 404

Digging Deeper Exercise 7 414

Chapter 8 Physical Optics 417

Miscellaneous Worked Examples 469

Digging Deeper Exercise 8 476

Chapter 9 Sound, Electromagnetic Waves, and the Doppler Effect 479

Miscellaneous Worked Examples 533

Digging Deeper Exercise 9 542

Examination Questions for Part Two 545

PART THREE: ELECTRICITY

Chapter 10 Electrostatics 552

Miscellaneous Worked Examples 625

Digging Deeper Exercise 10 637

Chapter 11 Capacitance 640

Miscellaneous Worked Examples 690

Digging Deeper Exercise 11 698

Chapter 12 Current Electricity 701

Miscellaneous Worked Examples 784

Digging Deeper Exercise 12 790

Chapter 13 Alternating Current Theory 795

Miscellaneous Worked Examples 838

Digging Deeper Exercise 13 844

Examination Questions for Part Three 847

Answers to Digging Deeper Exercises 855

Solutions to Examination Questions 896

Part One 896

Part Two 907

Part Three 917

Index 931

PREFACE

This book was written because you deserve better.

Advanced Level physics demands deep understanding, yet too many students meet heat, waves, and electricity as three disconnected lists of formulas to be memorised: a rule for expansion here, a rule for lenses there, a rule for circuits somewhere else, none of them explained and none of them connected. They memorise what they do not understand, pass examinations they cannot apply, and leave Form Six believing that physics is a collection of equations to be recalled under pressure. That is not physics. That is memorisation dressed as understanding.

Physics is the study of how the universe works. It is the reason a kettle boils, a radio carries a voice clear across the country, and a phone charger turns the violence of the mains into the gentle trickle of current that wakes your screen. It is powerful, deeply satisfying, and far more logical than its reputation suggests. But understanding it requires a different kind of textbook: one that explains why before how, that derives every formula instead of dropping it from the sky, and that tests your thinking, not just your memory. That is what this book sets out to be.

Why Heat, Waves and Electricity?

The first book in this series, Foundations of Physics, built the language of mechanics: force, energy, momentum, and motion. This book puts that language to work. Heat, waves, and electricity are not three unrelated subjects; they are the three great ways in which matter and energy interact, which is why this volume is called Physics Interactions. Heat is the restless motion of particles far too small to see. Waves are how energy travels when matter itself stays behind, from a plucked string to the light that reaches your eye. Electricity is the interaction of charges through the invisible fields that fill the space around them. Between them they explain almost every machine you will ever use, and they carry an enormous portion of the Advanced Level syllabus. The student who learns to hear the single set of ideas running beneath all three finds that this whole region of physics begins to speak with one voice.

How this book is organised

The book is divided into three parts.

Part One: Properties of Matter and Thermal Physics studies how matter behaves under load and under heat: the stretching of solids, the true meaning of temperature, the flow of heat from one body to another, the ceaseless motion of gas molecules, and the laws of thermodynamics that govern every engine and refrigerator.

Part Two: Vibrations and Waves covers oscillations and the waves they give rise to: simple harmonic motion, wave motion, physical optics, and sound, electromagnetic waves, and the Doppler effect.

Part Three: Electricity carries you from a single charge at rest to the alternating current in the wall: electrostatics, capacitance, current electricity, and alternating current theory.

Within each chapter, the content is built in layers. Concepts are developed through narrative explanations, followed by in-text worked examples in three tiers:

BINDER examples anchor the basics: if you can do these, you own the concept.

REAL examples connect physics to everyday Tanzanian life, showing that physics is not confined to the laboratory.

HOT examples push you to examination standard and beyond, combining several concepts at once and demanding strategy.

Each in-text example ends with **Making Sense of the Answer** (which checks whether the result is physically reasonable) and **Think Like a Physicist** (which teaches a transferable problem-solving habit you can carry to any future problem).

At the end of each chapter, **Miscellaneous Worked Examples** combine ideas from across the chapter in the way real examinations do, followed by a **Digging Deeper Exercise** with questions in the same three tiers (BINDER, REAL, HOT).

At the end of each part, **Examination Questions** modelled on NECTA, Mock, Regional, and Commonwealth papers test everything from that part.

Near the end of the book, you will find complete answers to every Digging Deeper question, including full explanations for the explanatory questions, not just numerical answers. Full worked solutions to all Examination Questions are provided as well. Nothing is left unanswered. If you are stuck, the solution is there to learn from, but only after you have attempted the problem honestly with the book closed.

Three people will accompany you through every chapter.

Kipanga is the student who volunteers first and thinks afterwards. He is confident, energetic, and wrong more often than he would like to admit. But he learns from every mistake, and his mistakes will teach you too, because they are the same mistakes real students make in real classrooms.

Kipute is the student who sits in the second row with her notebook already open. She is curious, analytical, and asks the questions that separate good students from great ones. When Kipute challenges an idea, pay attention: the answer usually contains something important.

Mr. Akilikubwa is the teacher who never raises his voice but somehow makes every student listen. His explanations use vivid analogies drawn from everyday Tanzanian life, and his patience is legendary. When he speaks, the concept lands and stays.

These three characters are not decorations. They carry the teaching. Their dialogues address the exact misconceptions, confusions, and "almost-right" ideas that real students struggle with every day.

The team behind this book

Enjoy and Master Advanced Physics is the product of the Ngaiza Education Hub (NEH). It was built by a team of experienced physics educators who share a common conviction: that Tanzanian students deserve world-class teaching materials written for their context, their syllabus, and their aspirations.

Writers: Mr. Amir Juma, Mr. Kelvin Mashu, Mr. Endrew Mariki, and Mr. Morris Mulokozi.

Editors: Dr. Aloyce Julian and Dr. Alex Muganyizi.

Designer: Mr. Frank Mwita.

I am grateful to every member of this team for their dedication, their expertise, and their refusal to settle for "good enough." I also thank the students whose hard questions and honest frustrations shaped every chapter. They reminded us, constantly, that a textbook which cannot be understood by the student it was written for has failed, no matter how correct its physics may be.

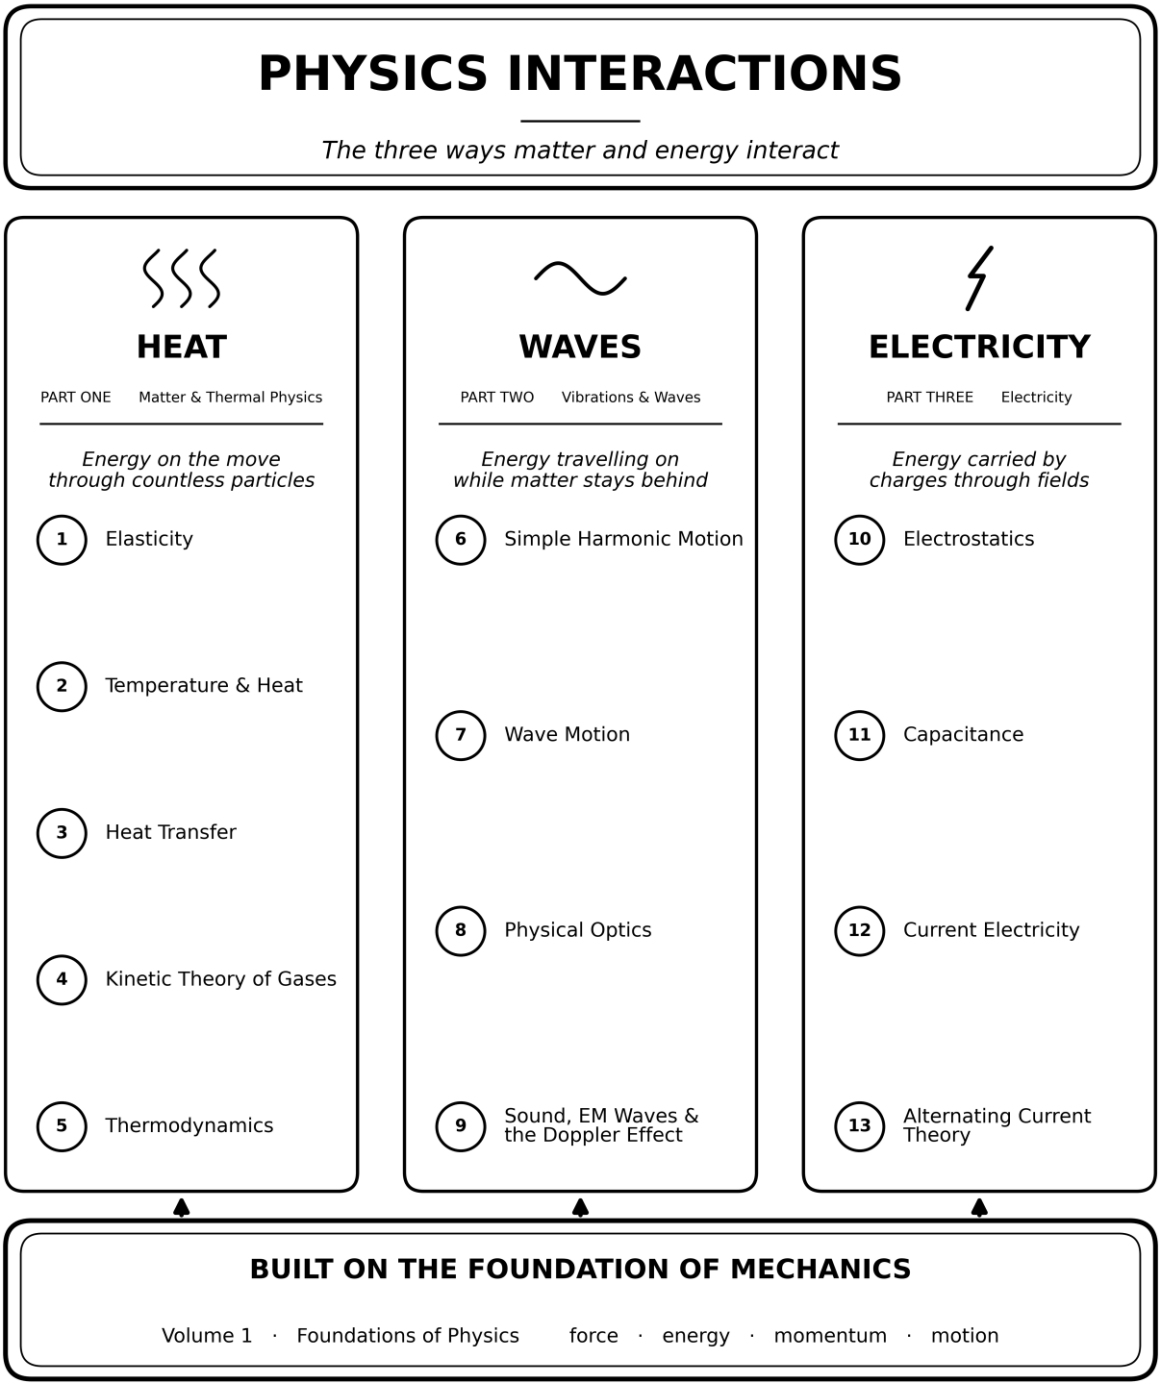
To every Advanced Level student opening this book: physics is not as hard as its reputation. It is as logical as a well-told story, as satisfying as a puzzle that finally clicks, and as useful as any skill you will ever learn. Give it your honest effort. This book will give you its honest best in return.

Ngaiza, Lusima
Series Creator and Director

This book is dedicated to every Tanzanian student who was told physics is too hard. It is not. And this book will prove it to you.

THE MAP OF THE INTERACTIONS

Everything in this book is one idea seen from three sides: how matter and energy interact. When a chapter feels unfamiliar, find where it sits on this map, and you will see at once what it connects to.



PART ONE

PROPERTIES OF MATTER AND THERMAL PHYSICS

Matter is not the silent backdrop against which physics happens. Matter has a structure, a temperament, and a set of habits of its own, and the Form Five physicist who learns to listen to those habits will find that an enormous portion of the syllabus suddenly speaks with a single voice. Steel bends and remembers, glass shatters, rubber stretches and warms, railway tracks lengthen in the sun, ice floats while stones sink, and a candle flame transfers heat across an empty room: every one of these is matter telling us something about its molecules. This part of the book is dedicated to learning that language. We will study how solids deform under load, how heat moves through substances and changes their state, and how the molecular picture quietly underwrites all of it.

Chapter 1

ELASTICITY

INTRODUCTION

Kipanga had taken many daladalas in his life, but this particular Tuesday afternoon would later become, in his own private filing system, the day a broken bench taught him advanced physics.

The route was Mwenge to Mbagala. The vehicle was the usual kind, with paint that had been refreshed so many times that you could no longer tell what colour it had started as. Kipanga boarded near the front, paid his fare to the conductor, and squeezed himself onto the long bench seat behind the driver, the one that runs the full width of the vehicle and is engineered, in theory, with a hinged cushion above two long springs.

In practice, only half the bench was still working. The right half was already occupied by a serious-looking woman with a basket of mangoes, who did not look like somebody in the mood to share bench physics. So Kipanga took the left half.

He sat. There was a quiet metallic sigh from beneath the cushion, and then nothing. The cushion did not move. It did not rise up to meet him, did not press back against him, did not behave in any way like a bench should behave. It was flat. Permanently, philosophically, unconditionally flat.

He bounced experimentally. The cushion did not bounce. Kipanga was sitting, essentially, on a metal pancake.

Across the aisle, a small boy of perhaps eight years old was on the bench opposite, swinging his legs and bouncing happily with every pothole the daladala hit, which was many. That bench, clearly, still had opinions. It compressed when the boy came down, and it pushed him back up again. Two benches. Same daladala. Presumably the same factory, the same kind of spring underneath. One was alive and one was dead.

The daladala arrived at Ubungo bus stop.

New passengers poured in. A man with a roll of foam mattresses squeezed past the conductor. A woman with a chicken in a box. Three secondary school students in uniform. Two office workers who had clearly missed the earlier bus. By the time the doors finally closed, every seat was full, the aisle was full, the space behind the driver was full, and several people were standing in positions the daladala engineers had probably never imagined.

One of those people was **Kipute.**

She found no seat, and ended up standing right next to Kipanga, one hand gripping the overhead rail, her school bag wedged awkwardly against the driver's seat, her physics notebook somehow still open in her free hand, because of course it was. From where she was standing, she could look straight down at the top of Kipanga's head.

"You look like a man with a question," she said.

Kipanga looked up. *"Kipute! When did you get on?"*

"Ubungo. About thirty seconds ago. You were staring at your bench as if it had insulted your mother."

"This bench is broken," said **Kipanga.** *"But that one isn't."* He pointed across the aisle at the bouncing boy.

"It is not broken," said Kipute. *"It has been overloaded."*

"What is the difference?"

"The difference is the chapter we start tomorrow."

"What chapter?"

"Elasticity. Mr. Akilikubwa has been threatening it for two weeks. He says it is the chapter where physics finally tells the truth about why some things spring back and others give up."

The daladala hit a pothole. Every standing passenger swayed as one body. Kipute braced herself against the overhead rail. Kipanga, sitting on his dead bench, felt absolutely nothing because a pancake cannot bounce.

"Why?" he said, with the kind of urgency only a bad bench can provide. *"Why does one bench remember its shape and the other one forgets? They are both made of metal. They were both pressed by the same kinds of bottoms. Why does one come back and the other one stay flat?"*

Kipute smiled, and closed her notebook, because she knew when an answer was going to be more useful if it came from Mr. Akilikubwa rather than from her.

"Tomorrow," she said, looking down at him from her standing position. *"Bring the question."*

———————

The next morning, Mr. Akilikubwa walked into Form Five Physics carrying two short metal springs. He had cut them from identical bicycle pumps the previous evening. One was neatly coiled and looked ready for work. The other had clearly been

stretched far beyond its limit and now hung from his fingers like a piece of dead rope. He placed both on the demonstration table and waited for the class to settle.

Kipanga's hand went up before the bell had finished ringing.

"Sir, I have a question from a daladala."

"The best questions usually are."

Kipanga described the two benches. The class laughed in the way Tanzanian students always laugh at daladala stories, because every one of them had a similar memory waiting somewhere in the back of their mind. Mr. Akilikubwa let the laughter run its full course; he was never in a hurry to interrupt laughter, and then he picked up the two springs and held them up side by side.

"Two springs," he said. *"Same metal. Same factory. Same length, originally. One of them I bent gently and let go, and it came back. The other I stretched until it could not return, and it stayed stretched. The metal did not break. Every iron atom that was in this spring yesterday is still in it today. And yet one of them is alive, and one of them is dead."*

He set the springs down.

"This is the chapter where physics finally tells you the truth about what is happening inside that metal. We are going to climb all the way down to the level of individual atoms, look at the forces between them, and from those forces we are going to derive every elastic thing in your syllabus. Hooke's law. Young's modulus. The reason a glass rod shatters and a steel wire bends. The reason a rubber band gets warm when you stretch and release it many times. The reason railway tracks have gaps. The reason an artery is stretchy and a tendon is not. All of it, from one picture."

He turned to the board and wrote four words in deliberate capitals: THE INTERMOLECULAR POTENTIAL ENERGY CURVE.

"By the end of this chapter, when somebody hands you a piece of dead spring on a daladala bench, you will not just shake your head and accept it as fate. You will know exactly which atoms moved past which other atoms, and why they are never coming back. Welcome to Chapter 1."

Kipanga, sitting in the front row, wrote a note in the margin of his exercise book that read: *"Daladala bench; atoms moved."* He underlined it twice.

INTERMOLECULAR FORCES AND THE POTENTIAL ENERGY CURVE

Mr. Akilikubwa's promise was bold, and he is about to make it good. Every elastic phenomenon in this chapter; every spring constant, every modulus, every thermal stress, every reason a daladala bench either remembers its shape or forgets it, flows from a single curve. We will draw that curve once, label it carefully, and then return to it again and again as the chapter unfolds. Every time some new elastic behaviour appears in this book, we will discover that it was already living, quietly, somewhere on this curve.

Forces between molecules

Two molecules in a solid do not just sit beside each other passively. They interact. Between every pair of neighbouring molecules there is a force that depends only on how far apart they are. This is the central physical fact on which everything else in this chapter is built, and it deserves to be stated slowly.

When the molecules are far apart, the force between them is attractive: each molecule pulls the other gently towards it. This is the force that holds a solid together at all. Without it, every block of metal in the world would simply fall apart into a cloud of atoms.

When the molecules are pushed too close together, however, the force becomes strongly repulsive: they shove each other away with great vigour. This repulsion is what stops a solid from being compressed indefinitely. It is the reason you cannot squeeze a steel ball into a smaller steel ball no matter how hard you press.

At one particular separation, and one only, the attractive and repulsive forces exactly cancel and the net force is zero. We give this separation a name and a symbol. We call it the **equilibrium separation**, and we write it r_0. It can be defined as:

The distance between two neighbouring molecules at which the attractive and repulsive intermolecular forces exactly cancel, so that the net force between them is zero and their potential energy is at its minimum.

A pair of bonded molecules will always settle, if left alone, at this separation. It is where they are most comfortable, and where they will return if disturbed gently.

From force to potential energy

It is often more useful to think about the energy of a pair of molecules than about the force directly, because energy is a single number while force has a direction. The two descriptions are completely equivalent: from one we can always get the other. The potential energy of a pair of molecules separated by a distance r, which we shall write U(r), is the work done by the intermolecular force as the molecules are brought from infinity to that separation. Since the molecules attract each other from

afar, the force does positive work on them as they approach, and the potential energy decreases. Since they repel each other strongly when very close, the potential energy rises sharply at small separations. Somewhere in between, the energy reaches its lowest possible value. That lowest value sits exactly at r_0, because that is where the force is zero and there is nothing left to push the energy down further.

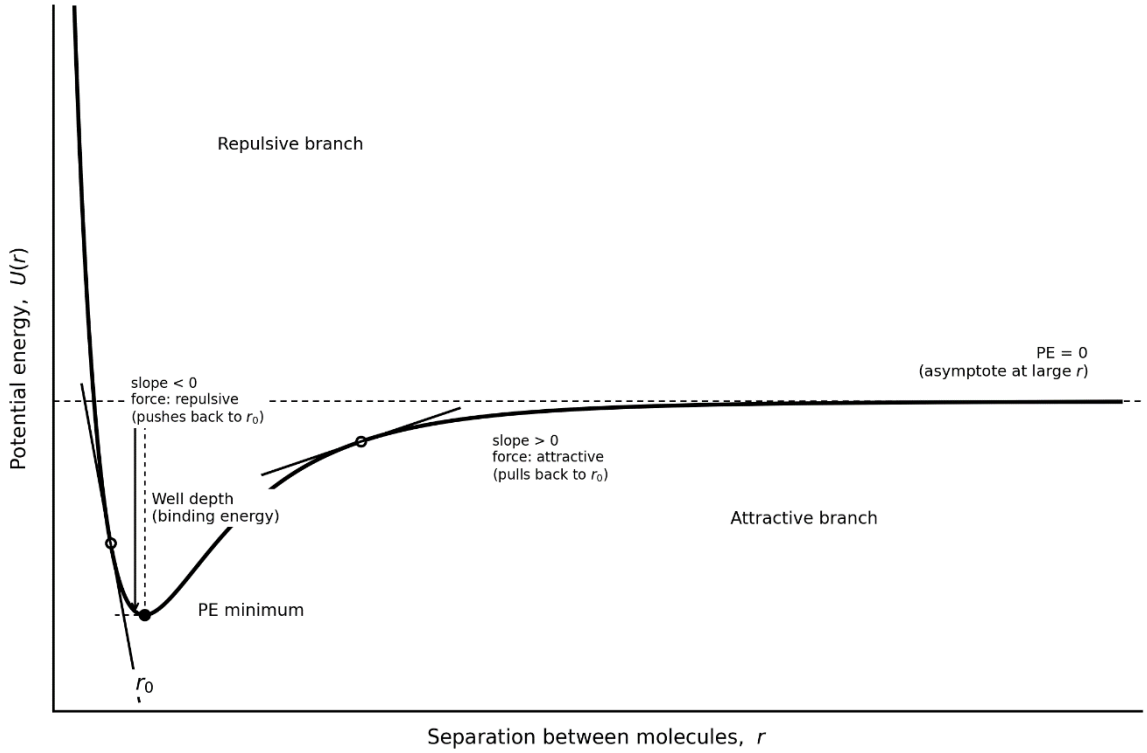


Figure: ***The intermolecular potential energy U(r) for a pair of bonded molecules as a function of their separation r.*** *The curve reaches its minimum at the equilibrium separation r_0, where the attractive and repulsive forces exactly balance and the net force is zero. The depth of the well below the horizontal asymptote at U = 0 is the binding energy of the pair. At separations smaller than r_0 the steep repulsive branch dominates, and at separations larger than r_0 the gentler attractive branch rises towards the asymptote, which represents two free molecules infinitely far apart. The slope of the curve at any point gives the magnitude and direction of the force on the molecules, and that force always points back towards r_0.*

Reading the curve

This curve is one of the most important pictures in the entire book. Every feature of it has a physical meaning, and we will now go through them one by one. Spend a few minutes with the diagram before reading on; almost everything that follows in this chapter is just a careful re-description of something already visible there.

The deep well at the centre. The curve has a clear minimum at the equilibrium separation. The depth of this well, measured from the dashed asymptote down to the bottom, is called the binding energy of the pair of molecules. It is the amount of energy you would have to supply to pull the two molecules apart from rest at equilibrium and carry them off to infinity. A deeper well means a more strongly bound pair of molecules, and therefore a more rigid material. A shallower well means a softer, more easily deformed material.

The repulsive branch on the left. To the left of the minimum, the curve climbs steeply and almost vertically. This is the regime in which the molecules are being pushed unwillingly into each other. Even a tiny decrease in separation costs an enormous amount of potential energy, because the slope of the curve is very steep here. A steep slope means a large force, and the direction of the force opposes the compression; the molecules push back hard. This is why solids resist compression so strongly: there is a wall of energy waiting on the left, and it gets steeper the further you try to climb.

The attractive branch on the right. To the right of the minimum, the curve rises again, but much more gently, and approaches the horizontal asymptote at PE = 0 from below. This is the regime in which the molecules are being pulled apart. Energy must be supplied to stretch them, but the slope of the curve is shallower here than on the left, so the restoring force is weaker for the same displacement. The curve eventually flattens out at large separations: once the molecules are far enough apart, they no longer feel each other at all. The asymptote represents two free, independent molecules that have completely escaped one another's influence.

Small displacements behave like a spring. Look very carefully at the bottom of the well, in a small region just to the left and right of the minimum. In that small region the curve looks almost like a parabola; a smooth, symmetrical U-shape. This is no accident. Any smooth curve, when viewed close enough to its minimum, looks like a parabola, and the potential energy

of a parabola corresponds to a force that is proportional to displacement. In other words: *for small displacements from equilibrium, the restoring force on a pair of molecules is directly proportional to how far they have been displaced.* This is exactly the statement of **Hooke's law**. We have just discovered the molecular origin of the most famous law in elasticity, and we have not written a single equation to do it. Later, we will give the law its proper algebraic form; for now it is enough to know that Hooke's law is not a separate principle imposed from outside. It is the bottom of the intermolecular well, viewed from close up.

Large displacements do not behave like a spring. Now look further out from the minimum, on either side. The curve is no longer a parabola. The left branch climbs faster than a parabola; the right branch climbs slower. The further you stretch or compress, the worse Hooke's law becomes as an approximation. This is why every spring in the world has a limit beyond which it stops behaving like a spring, and why Kipanga's daladala bench, having been pushed so far to the right of the well that it crested the slope and came down on the far side, never returned to its original shape.

The well is not symmetric. There is one more feature of the curve that will become very important later, on thermal expansion. Notice that the well is not symmetric: the repulsive wall on the left is much steeper than the attractive slope on the right. This asymmetry means that when a molecule vibrates back and forth in the well, it spends slightly more of its time on the right (the attractive side) than on the left. As the temperature rises and the vibrations become more energetic, this lopsidedness becomes more pronounced, and the average separation drifts to the right. The molecules move further apart on average, and the material expands. Thermal expansion, which has puzzled students for centuries, turns out to be nothing more than the asymmetry of the well, made visible. We will return to this in the thermal expansion section with the curve in front of us.

One curve, three classes of behaviour

The shape of this same curve, drawn for different materials, explains why those materials behave so differently. The picture is universal; the depth and width of the well are not. Three contrasting cases are worth carrying with you through the rest of the chapter.

1) A deep, narrow well (applies for a rigid solid such as steel or diamond)

The well is so deep that it takes an enormous amount of energy to pull the molecules apart, and so steep on both sides that even a small displacement from equilibrium produces a huge restoring force. Materials of this kind have very high binding energies, very high stiffness, and very high melting points. They resist both compression and extension fiercely.

2) A shallower, broader well (applies for a softer solid such as lead or copper)

The well is still a well, so the material still holds together, but the slopes are gentler. The same displacement from equilibrium produces a smaller restoring force. The material is more easily stretched, more easily compressed, more easily dented, and melts at a lower temperature.

3) No bound state at all (applies for a gas)

In a gas, the molecules are too far apart and moving too fast for the attractive part of the curve to capture them. They are not sitting at the bottom of any well; they are flying past each other out at the flat asymptote on the right. The gas as a whole feels essentially no restoring force when you try to stretch it, because there is nothing pulling its molecules back. This is exactly why a gas cannot resist being pulled apart: there is no spring to restore it. But push the molecules of a gas closer together, and they start to climb the repulsive wall on the left, and they push back. This is why a gas can be compressed against a force (think of pumping a bicycle tyre, even though it offers no resistance to expansion).

We now have the picture. Two molecules sit in a well, comfortable at the bottom, restored gently when nudged a little, restored more roughly when nudged further, and finally unable to return at all when nudged past the edge. Everything else in this chapter: the algebra of stress and strain, the three moduli, the energy stored in a stretched wire, the failure of an overloaded rail, the heat of a snapping rubber band, the thermal stress in a clamped composite rod, is the same picture, dressed in different clothes. Before we put any of those clothes on, however, we need a little practice reading the curve out loud. The two examples that follow are not calculations. They are exercises in seeing.

BINDER Example 1

Use the intermolecular potential energy curve to explain why a solid resists both compression and extension, while a gas resists only compression.

Solution

In a solid, the molecules sit close to the equilibrium separation r_0 at the bottom of the intermolecular potential energy well. Any attempt to compress the solid pushes neighbouring molecules to the left of the minimum, where the curve climbs the steep repulsive branch and produces a large restoring force pushing them apart again. Any attempt to stretch the solid pulls the molecules to the right of the minimum, onto the attractive branch, where the curve again rises and produces a restoring force pulling them back together. Either way, displacing the molecules from r_0 costs energy, and the energy gradient gives a restoring force. A solid therefore resists both compression and extension.

In a gas, the molecules are not sitting in the well at all. They are far out on the flat asymptote on the right, where the curve is essentially horizontal. Pulling them further apart produces no change in potential energy, and therefore no restoring force as there is nothing to pull the molecules of a gas back together once they have been separated. A gas offers no resistance to being stretched. Pushing the molecules of a gas closer together, however, eventually drives them down the asymptote and onto the rising part of the curve, where the slope is no longer zero and a restoring force appears. The gas resists compression.

The asymmetry of behaviour between solids and gases is therefore not a separate fact about each state of matter. It is one fact about the same curve, read in two different places.

Making Sense of the Answer: *The argument never invoked any special property of "solids" or "gases" as such. It used only the shape of the intermolecular potential energy curve and the rule that force is minus the slope of the energy. The same logic, applied at different points along the curve, gave the behaviour of two different states of matter. This is the kind of unification physicists live for.*

Think Like a Physicist: *Whenever a question asks you to compare the elastic behaviour of two materials or two states of matter, your first move should be to draw the curve and ask, "which part of the curve are the molecules sitting on?" Almost every qualitative question in this chapter can be answered by pointing at the right part of the curve. Reach for the picture before reaching for the words.*

REAL Example 2

Kipute is sitting in the front row, frowning at the curve on the board. After a long pause, she raises her hand.

Kipute: *"Sir, you said that molecules attract each other when they are far apart. But the air in this room is made of molecules, and they are far apart. Why doesn't all the air collapse into a puddle on the floor? If they attract each other, surely they should pull themselves into a tiny ball at my feet."*

Mr. Akilikubwa: *"That is exactly the question this curve was waiting for somebody to ask. The answer is on the diagram; you just have to look at it the right way."*

Use the intermolecular potential energy curve, together with the kinetic energy of the molecules, to explain why the air in the room does not collapse onto the floor under its own intermolecular attraction.

Solution

Two ingredients matter, and the answer is a comparison between them. The first ingredient is the depth of the well which reflects the binding energy that two air molecules would gain if they fell into the bottom of it. The second ingredient is the kinetic energy each molecule already has by virtue of being at room temperature, the energy of its restless thermal motion.

If the kinetic energy of an air molecule is much smaller than the depth of the well, then any two molecules that drift close enough to feel each other's attraction will be captured by it. They will fall into the well, lose their freedom, and form a bound pair. Many such pairs together make a liquid or a solid. This is exactly what happens to water vapour when it cools to a temperature at which kinetic energies are comparable to the depth of the water-water well: the vapour condenses into liquid water.

If the kinetic energy is much larger than the depth of the well, however, the situation is completely different. Two molecules that drift towards each other do feel the attractive part of the curve, but the attractive force is far too weak to capture them. They fall a little way into the well, gather a bit of extra speed on the way down, and then sail straight out the other side, still moving fast enough to escape. The well is still there. The attraction is still there. But the molecules have so much kinetic energy that the well looks, from their point of view, like a tiny dimple on the floor of an enormous room. They walk right over it without falling in.

At room temperature, the average kinetic energy of an air molecule is several times larger than the depth of the well between two air molecules. The attraction exists, but it is completely overpowered by the thermal motion. The air therefore stays as a gas, with its molecules flying past each other on the flat asymptote of the curve, never bound, never collapsing. Cool the air enough, say down to about −196 °C for nitrogen, and the kinetic energy will fall below the well depth, and at that point the air will indeed condense into a puddle on the floor. Liquid nitrogen is exactly that puddle.

Making Sense of the Answer: *The molecules of a gas are not free because they have no attraction. They are free because they have so much kinetic energy that the attraction is, for all practical purposes, invisible. Whenever you want to know whether a substance will be a solid, a liquid, or a gas, the question to ask is not "is there a force?" (there is always a force). The question to ask is "is the kinetic energy big enough to overcome (escape) it?"*

Think Like a Physicist: *Two energies, compared. Almost every question in physics about whether something will stay bound or escape; a satellite in orbit, an electron in an atom, a molecule in a liquid, a footballer trying to leave the pitch, reduces to the same comparison: the depth of the well versus the kinetic energy of the thing trying to escape it. Once you start noticing this pattern, you will see it in chapter after chapter for the rest of your physics life.*

The curve is now in our hands, and we can read it. The next section takes the bottom of the well and gives it the algebraic dress that engineers and examination boards prefer. We will define the two quantities: **stress** and **strain**, that physicists use

to describe loaded materials, and we will see that they are nothing more than careful, fair ways of measuring how far we have pushed the molecules from r_0, and how hard we had to push to do it.

STRESS AND STRAIN

At the bottom of the intermolecular potential energy well introduced in the previous section, the restoring force on a pair of molecules is proportional to how far they have been displaced from their equilibrium separation r_0. That is the molecular truth. We now have to dress that truth in the algebra of engineering, because in the laboratory we never load a single pair of molecules. We load whole wires, whole rods, whole rails, and whole bones, all of which contain countless molecular pairs being pulled or pushed at the same time. Two questions immediately arise. *First, how hard are we pushing? Second, how far have we gone?* The answers to these two questions are the two quantities at the heart of every elasticity problem you will ever solve. They are **stress** and **strain**.

Both quantities have to be defined in a way that is fair across wires of different thicknesses, different lengths, and different materials. In physics, the word "fair" almost always means "as a ratio", and we are about to see why.

A small note before we begin. Readers who paid careful attention to the chapter on dimensional analysis in *Foundations of Physics* will remember that strain was listed there as one of the dimensionless quantities, alongside angles in radians. That listing was a preview of what we are about to define. The same word "elastic" also appeared in *Foundations of Physics*, in the chapter on collisions, but its meaning there was different. An elastic collision is one in which kinetic energy is conserved, whereas in this chapter **elastic** refers to *the ability of a material to return to its original shape after a load is removed.* Same word, two different physical uses, and the context will always make clear which one is meant.

Stress: how hard the molecules are being pushed

Imagine two steel wires of the same material and the same length, hanging vertically from a beam. From each wire we hang a 20N load. One wire is thick, with about the diameter of a ballpoint pen refill. The other is thin, with about the diameter of a strand of hair. The same 20N pulls on both. But the load reaches a very different number of molecules in each case. In the thick wire, the 20N is shared between an enormous number of molecular bonds across a wide cross section, and each bond carries only a tiny share. In the thin wire, the same 20N is shared between far fewer bonds, and each one of them is being pulled much harder. The thin wire feels the load more intensely than the thick one, even though the total force is identical in the two cases.

The total force, on its own, is therefore not a good measure of how hard the molecules are being pushed. What we really want to measure is the force per molecule in the cross section, which is the same as the force per unit area of that cross section. We give this ratio a name and a symbol. We call it **tensile stress**, and we write it **σ:**

$$\boldsymbol{\sigma = \frac{F}{A}}$$

where F is the magnitude of the force pulling along the wire, and A is the cross-sectional area of the wire measured perpendicular to the direction of the force.

By definition, **tensile stress** *is the tensile force acting on a body divided by the cross-sectional area of the body measured perpendicular to that force.* Its symbol is σ and its SI unit is the pascal (Pa), where 1Pa is equal to 1 N/m².

The pascal is a small unit. A wire being loaded hard enough to be of any practical interest is almost always under a stress of millions of pascals, so the megapascal ($1\ \text{MPa} = 10^6\ \text{Pa}$) and the gigapascal ($1\ \text{GPa} = 10^9\ \text{Pa}$) appear in almost every numerical answer in this chapter. Always be ready to convert between pascals, megapascals, and gigapascals.

Stress is not a vector in the everyday sense, because it does not have a single direction, but it does have a sign convention.

- When the force is **pulling the material apart**, the stress is called tensile and is taken as **positive**.
- When the force is **squeezing the material together,** the stress is called compressive and is taken as **negative.**
  By definition; **compressive stress** *is the stress that arises when the applied force is pushing the material together rather than pulling it apart*

The same algebraic formula $\sigma = \frac{F}{A}$ applies to both, with F carrying the appropriate sign. In this chapter we shall be working mostly with tensile loads and shall write everything as positive, but the convention is worth holding on to for the day when a railway track or a roof beam is in compression.

Tensile stress and compressive stress, taken together, are sometimes called **linear stress** or **longitudinal stress**, because in both of them the force is directed along a single line perpendicular to the cross section of the body. Later in this chapter we shall meet two other kinds of stress that act in different geometric ways, namely shear stress, in which the force is parallel to the surface rather than perpendicular to it, and bulk stress, in which the force is a pressure squeezing the body equally from every direction at once. For now, every stress in our calculations will be of the linear kind, which is what most of the syllabus and most examination questions are concerned with.

Cross-sectional area of a circular wire

Almost every wire in this book has a circular cross section, and the area of a circle is needed in nearly every numerical example. The two equally useful forms are:

$$A = \pi r^2 = \frac{\pi d^2}{4}$$

where r is the radius and d is the diameter. The diameter form is the one to use whenever an examination question quotes a diameter directly, because it removes one division and one chance to make a mistake. The radius form is the one to use whenever the question quotes a radius. Always convert millimetres to metres before substituting, because the SI unit of stress is built on the metre and not on the millimetre, and a single forgotten conversion will produce an answer that is wrong by a factor of a million.

Strain: how far the molecules have moved

Now consider the second question. The 20N load has produced an extension. *How far has the wire actually stretched?* Suppose it has lengthened by half a millimetre. Is half a millimetre a lot, or is it a little? The answer depends entirely on how long the wire was to begin with. Half a millimetre on a wire that started two metres long is a feather of an extension and is almost nothing. Half a millimetre on a wire that started two millimetres long is a strain of one quarter, which is far beyond what any metal could survive, and is a catastrophe. The absolute extension on its own tells us very little. What we really want to know is the fractional extension, which is the extension expressed as a ratio of the original length. We give this ratio its own name and symbol. We call it **tensile strain**, and we write it **ε**:

$$\boldsymbol{\varepsilon = \frac{\Delta L}{L}}$$

where L is the original (unloaded) length of the wire and ΔL is the change in length produced by the load.

By definition; **tensile strain** *is the change in length of a body divided by its original length.* Its symbol is ε. Because it is a length divided by a length, strain is a pure number with no unit.

Strain is therefore the dimensionless quantity you have already met under that name in the dimensional analysis discussion mentioned earlier. A bare number like 0.00025 is a perfectly good strain. It is also common to express strain as a percentage by multiplying by a hundred, so the same value would be reported as 0.025%, which sometimes makes it easier to picture how much the wire has actually changed.

The same sign convention applies as for stress. When the wire has lengthened, ΔL is positive and the strain is called tensile. When the wire has shortened under compression, ΔL is negative and the strain is called compressive. Same formula, signed inputs.

So **compressive strain** *is the strain that arises when the change in length of a body is a decrease rather than an increase.*

Tensile strain and compressive strain share an important feature. In both cases, the change we are measuring is a change in length along a single direction, the direction of the applied force. For this reason, the two are grouped together under a single umbrella name. They are together called **linear strain**, because the deformation happens along one line.

So **linear strain** *is the change in length of a body, in a chosen direction, divided by its original length in that direction.* Both tensile strain and compressive strain are the two cases of linear strain, the first when the body lengthens and the second when the body shortens.

With the language of stress and strain in our hands, we are now ready to do some calculations. The three examples that follow are deliberately straightforward. They build the muscle memory of converting millimetres to metres, computing circular cross sections, dividing forces by areas, and recognising stress and strain when they appear in the wild.

BINDER Example 3

A steel wire of radius 0.5mm hangs from a fixed support and carries a load of 20N at its lower end. Calculate the tensile stress in the wire.

Solution

Converting the radius to metres:

$$r = 0.5\text{mm} = 0.5 \times 10^{-3}\,\text{m} = 5 \times 10^{-4}\,\text{m}$$

The cross-sectional area of the wire is:

$$A = \pi r^2 = \pi \times (5 \times 10^{-4}\,\text{m})^2 = 7.854 \times 10^{-7}\,\text{m}^2$$

The tensile stress is the force divided by this area:

$$\sigma = \frac{F}{A} = \frac{20\text{N}}{7.854 \times 10^{-7}\text{m}^2}$$

$$\sigma = 2.546 \times 10^{7}\,\text{Pa} \approx 25.5\,\text{MPa}$$

The tensile stress in the wire is approximately 25.5MPa.

Making Sense of the Answer: *Twenty-five megapascals sounds like an enormous number, but for steel it is comfortably small. The elastic limit of structural steel is about 250MPa, ten times this value, and the breaking stress is several times higher again. The wire is therefore nowhere near in danger, and is doing exactly the kind of routine work for which steel was invented. The fact that the stress is millions of pascals is not a sign of trouble. It is the natural scale on which engineering metals operate.*

Think Like a Physicist: *Whenever a question gives you a radius or a diameter in millimetres, your very first move is to convert it to metres. The radius enters the area as a square, so a single missing factor of a thousand becomes a missing factor of a million in the answer. Convert first, substitute second, and check the units of the final result.*

BINDER Example 4

A copper rod of original length 2.0m extends by 0.5mm when a tensile load is applied to it. Calculate the tensile strain in the rod and express it as a percentage.

Solution

Converting the extension to metres so that the units cancel cleanly:

$$\Delta L = 0.5\text{mm} = 0.5 \times 10^{-3}\text{m} = 5 \times 10^{-4}\text{m}$$

The tensile strain is the fractional extension:

$$\varepsilon = \frac{\Delta L}{L} = \frac{5 \times 10^{-4}\ \text{m}}{2.0\ \text{m}}$$

$$\varepsilon = 2.5 \times 10^{-4}$$

Expressed as a percentage:

$$\varepsilon\,(\%) = 2.5 \times 10^{-4} \times 100\% = 0.025\%$$

The strain in the rod is 2.5×10^{-4}, or 0.025%.

Making Sense of the Answer: *A strain of 0.025% means the rod has lengthened by twenty-five parts per hundred thousand. The metres in the numerator and the metres in the denominator have cancelled completely, leaving a bare number with no unit, just as the definition demands. Strains in the elastic range of common metals are almost always smaller than 0.1%, which is the source of one of the perpetual frustrations of experimental elasticity. The quantity you most want to measure is always smaller than the quantity you can comfortably see. A later part of this chapter, on the experimental determination of Young's modulus, will return to this problem and show how Searle's apparatus is built specifically to defeat it.*

Think Like a Physicist: *When you compute a strain, the very first thing to check is that the answer has no units. If you find a stray metre or millimetre attached to your strain, you have either forgotten to convert one of the lengths or used inconsistent units. Strain is, by definition, a pure number, and that unit-cancellation is your first line of defence against arithmetic mistakes.*

REAL Example 5

A piano tuner stretches three steel strings between the same two pegs of a piano frame. The three strings are made of identical steel and are tightened to exactly the same tension, but their thicknesses are different. One is thick, one is medium, and one is thin. As the tuner raises the tension, the thinnest string is always the first to snap. Explain, using the definition of stress, why the thinnest string fails first even though all three strings carry the same applied force.

Solution

Steel breaks when the stress inside it reaches a definite value called the breaking stress. This breaking stress is a property of steel itself and not of any particular string. Every piece of the same kind of steel will fail at the same value of σ. The question is therefore not which string carries the largest force, because all three strings carry the same force by the conditions of the problem, but which string reaches the breaking stress first.

From the definition of stress,

$$\sigma = \frac{F}{A}$$

we see that for a fixed force F, the stress depends on the cross-sectional area A. A thick string has a large area, so the same force is shared between many molecules and the stress on each molecule is small. A thin string has a small area, so the same force is shared between far fewer molecules and the stress on each molecule is large. The thinner the string, the larger the stress, even though the total force pulling it has not changed at all.

As the tuner gradually increases the tension, the stress in all three strings rises together. But it rises fastest in the thin string, because the thin string starts at the highest stress and adds the largest stress increment for any given increase in tension. The thin string therefore reaches the breaking stress of steel first, while the medium and thick strings are still well below it. The thin string snaps, the tuner sighs, and physics smiles!

Making Sense of the Answer: *Notice that the answer never required a number. It rested entirely on the algebraic shape of the equation $\sigma = \frac{F}{A}$. When the same force is divided by a smaller area, the result is a larger stress. The physics of why thin*

things break first is exactly the physics of why high heels damage soft floors more than flat shoes do, why a sharp knife cuts more easily than a blunt one, and why an elephant's footprint sinks less into the ground than a stiletto's does. In every case the trick is the same. The same force, concentrated onto a smaller area, becomes a much larger stress.

Think Like a Physicist: *Whenever a problem describes "the same force" applied to objects of different shapes, your first instinct should be to compare not the forces but the stresses. The force is what is given to you, and the stress is what the molecules actually feel. Examination questions love to confuse the two, and students who confuse them lose marks they should not be losing.*

Stress and strain are now defined and we can compute them. But we still cannot connect them to one another. Knowing the stress in a wire does not yet tell us the strain it will produce, because we have not yet asked the wire how stiff it is. The relationship between stress and strain is governed by a single number that captures the elastic personality of the material, and that number is the subject of the next part of this chapter. We are about to meet Young's modulus, the most quoted material constant in all of engineering.

YOUNG'S MODULUS

Stress and strain are now sitting on the table in front of us, and we know how to calculate each of them. But we still cannot connect them to one another. *If we apply a certain stress to a wire, what strain will we get in return?* The answer depends on what the wire is made of. A steel wire and a rubber band subjected to the same stress will produce wildly different strains, and the reason is nothing more than the difference in stiffness between steel and rubber at the molecular level. To capture that stiffness in a single number, one quantity for each material, we need a new idea. That idea was first written down carefully by the English physicist Thomas **Young** in the year 1807, and it has carried his name ever since.

From the PE curve to a single number

The bottom of the intermolecular potential energy well is a parabola, as we saw in the section on the PE curve. A parabolic well gives a restoring force that is proportional to displacement, which is Hooke's law at the molecular level. What Young's modulus does is translate that molecular fact into the language of stress and strain, so that we can use it for whole wires and whole rods without having to count individual molecules.

The translation is disarmingly simple. *For any material being loaded within the elastic range, the stress is directly proportional to the strain.* That is **Hooke's law**, restated in the new language. The constant of proportionality is what we call **Young's modulus**, and we write it **E** (sometimes denoted by Y in older texts):

$$\boldsymbol{\sigma = E\varepsilon}$$

Equivalently, and more usefully for calculations, Young's modulus is the ratio of stress to strain:

$$\mathbf{E = \frac{\sigma}{\varepsilon}}$$

By definition; **Young's modulus** of a material *is the ratio of the tensile stress acting on the material to the resulting tensile strain, provided the material is within its elastic range.* Its symbol is E and its SI unit is the pascal (Pa), the same as stress, because strain is dimensionless.

Two consequences of this definition are worth stating immediately, because they are the single largest source of confusion in the whole topic.

First, Young's modulus is a property of the material and not of the particular wire or rod that has been cut from it. A thick steel cable and a thin steel hair have the same Young's modulus, because they are made of the same steel. Their stiffnesses as objects are very different, because one is thick and one is thin, but the stiffness per unit of material is the same. This is exactly the property we were hunting for in the previous section when we insisted that stress and strain be defined as ratios rather than as raw forces and raw extensions. Young's modulus is the reward for that discipline.

Second, the SI unit of Young's modulus is the same as the SI unit of stress, because strain has no unit at all. A pascal of stress divided by a dimensionless strain is still a pascal. In practice, the pascal is far too small for engineering materials. Steel has a Young's modulus of roughly two hundred billion pascals, which is why nearly every number you will see in this section is quoted in gigapascals, where 1GPa is equal to one billion pascals, or 10^9Pa.

The expanded form for laboratory use

The definition $E = \frac{\sigma}{\varepsilon}$ is compact and elegant, but in the laboratory we almost never measure stress and strain directly. We measure force, area, original length, and extension. Substituting $\sigma = \frac{F}{A}$ and $\varepsilon = \frac{\Delta L}{L}$ into the definition gives the expanded form:

$$\mathbf{E = \frac{\frac{F}{A}}{\frac{\Delta L}{L}} = \frac{FL}{A\Delta L}}$$

Every one of its four inputs (force, original length, cross-sectional area, and extension) is something you can read off a laboratory bench.

The same equation can be rearranged to give each of its variables in turn. The four useful rearrangements are:

$$F = \frac{EA\Delta L}{L}$$

$$\Delta L = \frac{FL}{EA}$$

$$A = \frac{FL}{E\Delta L}$$

$$L = \frac{EA\Delta L}{F}$$

You will use the first two of these constantly. The first one tells you the force needed to produce a chosen extension in a wire of known material and known dimensions. The second tells you the extension produced in a wire of known material and known dimensions by a given force. Between them, they solve most of the numerical elasticity problems you will ever meet.

Young's modulus values for common materials

Young's modulus values span an enormous range, from gigapascals for the hardest engineering metals down to megapascals for rubber and soft tissue, a factor of roughly one hundred thousand from the stiffest material in the table to the softest. The table below lists values you will meet throughout this chapter. You are not expected to memorise the numbers, but you should develop a feel for their order of magnitude, because a great many examination questions can be sanity-checked simply by noticing whether the answer you computed sits in the right neighbourhood for the material concerned.

Material	**Young's modulus E**
Diamond	1100GPa
Steel	200GPa
Copper	120GPa
Brass	100GPa
Aluminium	70GPa
Glass	70GPa
Bone (compact)	20GPa
Concrete	25GPa
Wood (along grain)	10GPa
Tendon	1.5GPa
Rubber	0.005 to 0.1GPa

Three patterns in the table are worth noticing before we move on.

1) The hardest engineering metals (steel, copper, brass, aluminium) all sit between 70 and 200 gigapascals, and diamond at the top of the list is an outlier with a value several times higher.
2) Among the biological tissues, bone is remarkably stiff for a living material, sitting at 20 gigapascals, which places it in the same general neighbourhood as ordinary structural concrete. This is one of the reasons the human skeleton can support the body against gravity for decades. Tendon is softer than bone by a factor of more than ten, which is exactly what you want from a tissue whose job is to stretch slightly under load and store elastic energy.
3) Rubber is softer still, by another factor that runs from roughly fifteen for the stiffer rubbers up to several hundred for the softest ones, which is why a rubber band can be stretched to several times its natural length without breaking, something no metal in the table could ever do.

From Young's modulus to the spring constant of a wire

Before we can compute the spring constant of a wire, we need to introduce two ideas that have been quietly waiting in the background of the chapter. The first is the spring constant of an ordinary spring. The second is the law that governs how that spring behaves under load, which is called **Hooke's law.**

Take an ordinary helical spring of the kind found inside a daladala bench, a kitchen weighing scale, or a retractable ballpoint pen. Hang it vertically from a fixed support, attach a small load to its lower end, and measure how much the spring stretches.

Now double the load and measure again. The extension also doubles. Triple the load, and the extension triples. Provided you do not load the spring so heavily that it deforms permanently, the extension is always proportional to the applied force. This proportionality is one of the oldest experimental laws in physics. It was first published by the English natural philosopher Robert Hooke in 1678. The law states that:

Provided a spring is not loaded beyond its elastic limit, the extension of the spring from its natural length is directly proportional to the applied force.

Mathematically:

$$\mathbf{F = kx}$$

where F is the force applied to the spring, x is the extension of the spring measured from its natural unloaded length, and k is a positive constant called the **spring constant**.

So the spring constant of a spring *is the constant of proportionality between the applied force and the extension in Hooke's law.* Its SI unit is newtons per metre (N/m). It tells you how much force is needed to stretch the spring by one unit of extension. It is a property of the particular spring, depending on the material it is made of and on the way it has been wound, and not a property of the load placed on it. A stiff spring has a large spring constant, because a large force is needed to produce even a small extension. A soft spring has a small spring constant, because a small force already produces a large extension.

Hooke's law is not a fundamental law of physics in the way Newton's laws are. It is an empirical regularity that holds very well for most materials over a small range of extensions and breaks down for large extensions, large compressions, and any loading that permanently deforms the spring. Inside that small range, however, it is one of the most useful relationships in all of physics, and the spring constant is the single number that captures it for any particular spring.

A wire loaded along its length behaves in exactly the same way as a spring, because a wire is just a very thin, very straight spring made of a single continuous piece of metal. It has a spring constant of its own, and we can compute that spring constant from the Young's modulus of the material it is made of, together with its length and its cross-sectional area.

Starting from the expanded form $E = \frac{FL}{A\Delta L}$, we rearrange for F:

$$F = \frac{EA\Delta L}{L}$$

$$F = \left(\frac{EA}{L}\right)\Delta L$$

Comparing this with Hooke's law in the form $F = k\Delta L$ (where the extension of the wire plays the role of x), we can read off the spring constant of the wire directly. The constant sitting in front of the extension is the wire's own spring constant:

$$\mathbf{k = \frac{EA}{L}}$$

So the spring constant of a wire loaded along its length *is the product of its Young's modulus and its cross-sectional area, divided by its original length.* It tells you how much force is needed to produce a unit extension in that particular wire.

Three features of this formula deserve to be stated out loud:

1) Making the wire longer reduces its spring constant $\left(k \propto \frac{1}{L}\right)$. A 2m wire has half the spring constant of an otherwise identical 1m wire made of the same material. Intuitively this is because the longer wire has more molecules in series between the anchor and the load, and each pair contributes a tiny extension of its own, so a longer stack of such pairs produces a larger total extension for the same force.
2) Making the wire thicker increases its spring constant $(k \propto A)$. A wire with twice the cross-sectional area has twice the spring constant. Intuitively this is because the thicker wire has more molecular bonds in parallel across its cross section, and each one carries a share of the load, so the load each bond has to carry is smaller and the extension each bond contributes is smaller.
3) Young's modulus itself is unchanged by any of this. Cutting a steel wire in half, or splitting it into two parallel halves, does nothing to the stiffness of the steel itself. The spring constants of the resulting pieces are different from the original, because the dimensions have changed, but E is still exactly what it was before. This is the same material-property-versus-object-property distinction we drew when we first defined Young's modulus as the ratio of stress to strain, showing up again in a new place. *The spring constant of a wire is an object property, but the Young's modulus of the metal the wire is made of is a material property,* and they should never be confused.

Stiffness is not the same as strength

Young's modulus measures stiffness, which is a material's resistance to being deformed. It does not measure strength, which is a material's resistance to being broken. These two ideas are constantly confused by students.

The cleanest way to see the distinction is with the Kipanga dialogue that follows, in which Kipanga finally asks the question that has been bothering him since he first saw a glass rod break cleanly on the demonstration table. The dialogue will not stop at the simple stiffness-versus-strength contrast, however. It will push one step further, into a third quantity called toughness,

because the everyday experience of dropping a glass bottle and watching it shatter is not really a question about breaking stress at all. It is a question about how much energy a material can absorb before its molecular bonds give way, and that is a slightly more subtle thing than either stiffness or strength on its own.

The stiffness of a material or of an object *is its resistance to elastic deformation under load. For a material, stiffness is measured by Young's modulus E. For a wire or rod, stiffness is measured by its spring constant k.*

The strength of a material, also called its **breaking stress**, *is the maximum stress the material can sustain before it fractures.* Its SI unit is the pascal, the same as stress itself. Stiffness and strength are two independent numbers, and a material can rank high on one scale and low on the other.

We shall say much more about strength when we reach the stress-strain curve later in this chapter and see how a material's elastic region, plastic region, and ultimate tensile strength all live on the same graph. For now, the important point is simply that stiffness and strength are two different questions with two different answers, and Young's modulus measures only the first of them.

With the theory now in place, we are ready for the first real numerical work of the chapter.

BINDER Example 6

A steel wire of length 3m and radius 0.6mm hangs from a ceiling and carries a load of 50N at its lower end. Young's modulus for steel is 2×10^{11}Pa. Calculate the extension of the wire.

Solution
Converting the radius to metres:

$$r = 0.6\text{mm} = 6 \times 10^{-4}\text{m}$$

The cross-sectional area of the wire is:

$$A = \pi r^2 = \pi \times (6 \times 10^{-4}\text{m})^2 = 1.131 \times 10^{-6}\ \text{m}^2$$

Using the rearrangement of Young's modulus that gives the extension:

$$\Delta L = \frac{FL}{EA}$$

$$\Delta L = \frac{50\text{N} \times 3\text{m}}{(2 \times 10^{11}\text{Pa})(1.131 \times 10^{-6}\ \text{m}^2)} = 6.63 \times 10^{-4}\text{m} \approx 0.66\text{mm}$$

The wire extends by about 0.66mm under the 50N load.

Making Sense of the Answer: *A 3m steel wire stretches by about two thirds of a millimetre under a 50N load. That is an extension you would struggle to see with the naked eye, and yet the wire is not doing anything remarkable. It is behaving exactly the way any steel wire of those dimensions would behave under a load of that size. This is why the elastic behaviour of metals is almost invisible in everyday life. The strains are small, the extensions are tiny, and the whole subject would be much harder to study if the deformations were not so well hidden by their own modest size.*

Think Like a Physicist: *The extension formula has four inputs. Any question that gives you any three of them will let you solve for the fourth. The fastest way to spot which rearrangement to use is to ask which variable is missing from the question, because that is always the variable you need to isolate. In this example the missing variable was the extension, so we used ΔL = FL/(EA). When the missing variable is F, the rearrangement is F = EAΔL/L. When the missing variable is E, the rearrangement is E = FL/(AΔL). Memorise the expanded form once and you get all four rearrangements for free.*

BINDER Example 7

A copper cable of diameter 4mm and length 5m extends by 2.1mm when a load F is applied to it. Young's modulus for copper is 1.2×10^{11}Pa. Calculate the load F.

Solution
Convert the diameter to metres and compute the cross-sectional area:

$$d = 4\text{mm} = 4 \times 10^{-3}\text{m}$$

$$A = \frac{\pi d^2}{4} = \frac{\pi \times (4 \times 10^{-3}\text{m})^2}{4} = 1.257 \times 10^{-5}\ \text{m}^2$$

Using the rearrangement that gives the force:

$$F = \frac{EA\Delta L}{L}$$

$$F = \frac{(1.2 \times 10^{11}\text{Pa})(1.257 \times 10^{-5}\ \text{m}^2)(2.1 \times 10^{-3}\text{m})}{5\text{m}} = 633.5\text{N} \approx 634\text{N}$$

The load on the cable is approximately 634N.

Making Sense of the Answer: *Six hundred and thirty-four newtons is roughly the weight of a 65kg adult standing on the cable. A 5m copper cable of 4mm diameter can easily support that weight while stretching by only 2.1mm, which is less than the thickness of a fingernail. Copper is softer than steel (its Young's modulus is about 60 percent that of steel), so the same cable made of steel would have stretched even less.*

Think Like a Physicist: *Whenever a question gives you a diameter rather than a radius, use the diameter form of the area formula directly. Squaring a diameter and dividing by four is one operation. Halving the diameter, then squaring, then multiplying by π is three operations, which is three chances to make a small arithmetic mistake. Fewer operations means fewer errors.*

REAL Example 8

Kipute is given two wires of unknown material in the physics laboratory. The first wire is 2m long with a diameter of 1mm. When she hangs an 80N load from it, the wire extends by 1.02mm. The second wire is 3m long with a diameter of 0.8mm. When she hangs a 50N load from this second wire, it extends by 2.49mm. Which of the two wires is stiffer, and by what factor?

Solution

The two wires differ in length, diameter, and applied load, so we cannot compare their stiffness directly from the extensions. We must compute Young's modulus for each wire and then compare the two values.

For the first wire, the cross-sectional area is:

$$A_1 = \frac{\pi d^2}{4} = \frac{\pi \times (1 \times 10^{-3}\text{m})^2}{4} = 7.854 \times 10^{-7}\ \text{m}^2$$

Young's modulus of the first wire:

$$E_1 = \frac{FL}{A_1 \Delta L_1}$$

$$E_1 = \frac{(80\text{N})(2\text{m})}{(7.854 \times 10^{-7}\ \text{m}^2)(1.02 \times 10^{-3}\text{m})} = 2.00 \times 10^{11}\text{Pa} \approx 200\text{GPa}$$

For the second wire, the cross-sectional area is:

$$A_2 = \frac{\pi \times (0.8 \times 10^{-3}\text{m})^2}{4} = 5.03 \times 10^{-7}\ \text{m}^2$$

Young's modulus of the second wire:

$$E_2 = \frac{(50\text{N})(3\text{m})}{(5.03 \times 10^{-7}\ \text{m}^2)(2.49 \times 10^{-3}\text{m})} = 1.20 \times 10^{11}\text{Pa} \approx 120\text{GPa}$$

The first wire is stiffer because its Young's modulus is higher. The ratio of the two moduli is:

$$\frac{E_1}{E_2} = \frac{200\text{GPa}}{120\text{GPa}} \approx 1.67$$

The first wire is about 1.67 times stiffer than the second.

Making Sense of the Answer: *The two wires differ in length, in thickness, and in the load they carry, so a direct comparison of their extensions would be meaningless. Only Young's modulus can strip away the geometry and leave a number that belongs to the material itself. The stiffer wire is close to steel at 200GPa and the softer wire is close to copper at 120GPa, which the student can confirm by glancing at the table earlier in this section.*

Think Like a Physicist: *Young's modulus exists precisely so that wires of different shapes loaded by different forces can still be compared fairly. Any question that asks which of several wires is stiffer, when the wires differ in more than one dimension or load, must be answered by computing Young's modulus for each wire separately. The raw extensions are meaningless on their own because they mix material properties with geometry.*

HOT Example 9

A 60N load hangs from a combination of two wires joined end to end. The upper wire is copper, with length 2m and cross-sectional area $1.5 \times 10^{-6}\ \text{m}^2$. The lower wire is steel, with length 2m and cross-sectional area $1 \times 10^{-6}\ \text{m}^2$. Young's modulus for copper is $1.2 \times 10^{11}\text{Pa}$, and for steel is $2 \times 10^{11}\text{Pa}$. Find (a) the tensile stress in each wire, (b) the extension of each wire, and (c) the total extension of the combination.

Solution

When two wires are joined end to end and a single load hangs from the lower wire, the same force is transmitted through both wires. Each wire therefore carries the full 60N, just as every section of a single long rope carries the full weight hanging from it. This is the key physical insight for the problem, and it is the one students most often miss.

(a) Stresses in the two wires:
For the copper wire:

$$\sigma_{Cu} = \frac{F}{A_{Cu}} = \frac{60\text{N}}{1.5 \times 10^{-6}\ \text{m}^2} = 4 \times 10^{7}\text{Pa} = 40\text{MPa}$$

For the steel wire:

$$\sigma_{St} = \frac{F}{A_{St}} = \frac{60\text{N}}{1 \times 10^{-6}\ \text{m}^2} = 6 \times 10^{7}\text{Pa} = 60\text{MPa}$$

The stresses are different even though the force is the same, because the cross-sectional areas are different. This is exactly the piano-string logic in the REAL Example 5.
(b) Extensions of the two wires:
For the copper wire, using the extension form of Young's modulus:

$$\Delta L_{Cu} = \frac{FL}{E_{Cu}A_{Cu}}$$

$$\Delta L_{Cu} = \frac{(60N)(2m)}{(1.2 \times 10^{11}Pa)(1.5 \times 10^{-6}\,m^2)}$$

$$\Delta L_{Cu} = 6.67 \times 10^{-4}m \approx 0.67mm$$

For the steel wire:

$$\Delta L_{St} = \frac{(60N)(2m)}{(2 \times 10^{11}Pa)(1 \times 10^{-6}\,m^2)}$$

$$\Delta L_{St} = 6 \times 10^{-4}m = 0.6mm$$

(c) Total extension
The wires are joined end to end, so their extensions add directly:

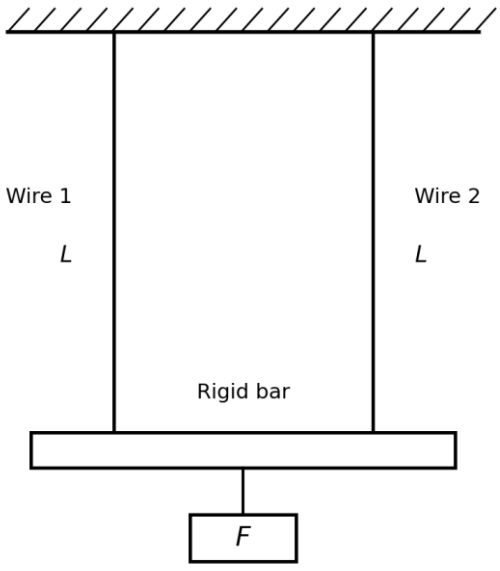
$$\Delta L_{total} = \Delta L_{Cu} + \Delta L_{St} = 0.67mm + 0.6mm = 1.27mm$$

The combined wire extends by approximately 1.27mm under the 60N load.
Making Sense of the Answer: *The copper wire stretches slightly more than the steel wire (0.67mm versus 0.6mm), even though the copper is thicker and carries a smaller stress. The reason is that copper's Young's modulus is only about 60 percent of steel's, and that softness more than compensates for the larger cross section. This is a good reminder that in composite problems you cannot predict which piece will stretch more just by looking at the dimensions. You have to do the arithmetic with both the dimensions and the Young's modulus of each material, because both contribute to the answer.*

Think Like a Physicist: *In any series combination of wires or springs, the load on each element is the same (the whole load passes through each piece one after the other), and the total extension is the sum of the individual extensions. In any parallel combination of wires or springs, the extension is the same (both pieces are forced to stretch by the same amount because they share the same two endpoints), and the total load is the sum of the individual loads. These two rules between them handle nearly every composite problem in elasticity, and they are the exact twins of the rules for resistors in series and parallel that you will meet later when you study electric circuits.*

HOT Example 10

A rigid horizontal bar is suspended from a fixed ceiling by two vertical wires attached to its two ends. Both wires have the same length of 2m. The first wire is steel with cross-sectional area 2×10^{-6} m^2 and Young's modulus 200GPa. The second wire is copper with cross-sectional area 3×10^{-6} m^2 and Young's modulus 120GPa. A 100N load is hung from the exact centre of the bar. If the bar is constrained to remain horizontal (i.e. it cannot rotate); find: (a) the extension of each wire, (b) the tension in each wire, and (c) confirm that the two tensions add up to the applied load.



Solution

Since the bar is **constrained to remain horizontal**, the two ends of the bar must move downward by exactly the same distance when the load is applied. Consequently, both wires undergo the same extension, making this a **parallel** arrangement wires.

The spring constant of each wire, from the formula $k = \frac{EA}{L}$, is:

$$k_1 = \frac{(200 \times 10^9 \text{Pa})(2 \times 10^{-6}\ \text{m}^2)}{2\text{m}} = 2 \times 10^5 \text{N/m}$$

$$k_2 = \frac{(120 \times 10^9 \text{Pa})(3 \times 10^{-6}\ \text{m}^2)}{2\text{m}} = 1.8 \times 10^5 \text{N/m}$$

(a) Extension of each wire (common extension):

In a parallel arrangement, the two spring constants add to give the effective stiffness of the combination:

$$k_{\text{total}} = k_1 + k_2 = (2 + 1.8) \times 10^5 \text{N/m} = 3.8 \times 10^5 \text{N/m}$$

Using $F = kx$ or $x = \frac{F}{k}$, the common extension is then:

$$x = \frac{F}{k_{\text{total}}} = \frac{100\text{N}}{3.8 \times 10^5 \text{N/m}} = 2.63 \times 10^{-4}\text{m} \approx 0.263\text{mm}$$

The extension of each wire is 0.263mm.

(b) Tension in each wire:

Applying Hooke's law F = kx to each wire separately:

$$F_1 = k_1 x = (2 \times 10^5 \text{N/m})(2.63 \times 10^{-4}\text{m}) = 52.6\text{N}$$

$$F_2 = k_2 x = (1.8 \times 10^5 \text{N/m})(2.63 \times 10^{-4}\text{m}) = 47.4\text{N}$$

The steel wire carries a tension of about 52.6N and the copper wire carries a tension of about 47.4N.

(c) Check that the tensions add up to the load.

$$F_1 + F_2 = 52.6\text{N} + 47.4\text{N} = 100\text{N}$$

The sum of the two tensions equals the applied load, as expected.

Making Sense of the Answer: *The steel wire takes the larger share of the load, about 53 percent of it, even though both wires have stretched by exactly the same amount. The reason is that steel is stiffer than copper, and a stiffer wire requires a larger force to produce the same extension. This is the opposite pattern from the series case we just met, where the softer wire stretched more and the stresses were different. In the parallel case, the extensions are forced to be the same and it is the forces that differ. The stiffer element always pulls its weight a bit harder than its softer neighbour, which is why the load-bearing cables in real structures are almost always made of the stiffest practical material the engineer can afford.*

Think Like a Physicist: *The spring constant of the parallel combination is the sum of the individual spring constants, which is the same addition rule as for conductances in parallel electrical circuits. The series case had extensions adding rather than spring constants adding, which corresponds to resistances adding in series electrical circuits. These parallels between elasticity and electricity are not a coincidence. Both are linear-response theories, where a driving quantity (force or voltage) produces a response (extension or current) proportional to itself, and the combination rules for series and parallel elements follow from the same underlying linearity regardless of the physical context.*

Young's modulus is now in your hands. You can compute it from raw laboratory measurements, you can use it to predict the extension of a wire under load, you can use it to compute the spring constant of a wire from its dimensions, and you can tell the difference between stiffness and strength. But so far we have stayed inside the elastic region, where loading and unloading a material returns it to its original shape with no memory of what happened in between. Real materials do not stay elastic forever. Push them hard enough and they begin to deform permanently. Push them harder still and they break. The full journey from gentle loading to final fracture is recorded on a single graph that every engineer in the world has committed to memory, and that graph is the subject of the next part of this chapter. Welcome, in advance, to the stress-strain curve.

THE STRESS-STRAIN CURVE

Young's modulus tells us how much a material deforms under a small load, provided the load is small enough for Hooke's law to apply. But Hooke's law does not apply forever. Pull on a metal wire hard enough and it will eventually stop being linear. Pull harder still and it will begin to deform permanently, so that when you release the load it does not return to its original length. Pull harder still and it will narrow locally, then snap. All of this behaviour, from the gentlest initial stretch to the final fracture, is captured on a single graph that every engineer in the world has committed to memory. The graph is called the stress-strain curve, and it is the subject of this section.

The curve is plotted with stress on the vertical axis and strain on the horizontal axis, using the definitions of both quantities that we established earlier in this chapter. To build the curve experimentally, you take a standard specimen of the material you want to test, you load it in tension along its length, you measure the extension as the load increases, and you compute the stress and strain at each measurement. Plotting the results gives a curve whose shape depends only on the material, not on the specimen dimensions. This is the same intensive-versus-extensive principle that gave us Young's modulus, applied now across the full range of loading.

The journey along a ductile stress-strain curve

We shall walk through the curve point by point, starting at the origin and ending at fracture. The figure below shows a typical stress-strain curve for a ductile material such as mild steel, with six key points labelled O, A, B, C, D, and E along the way.

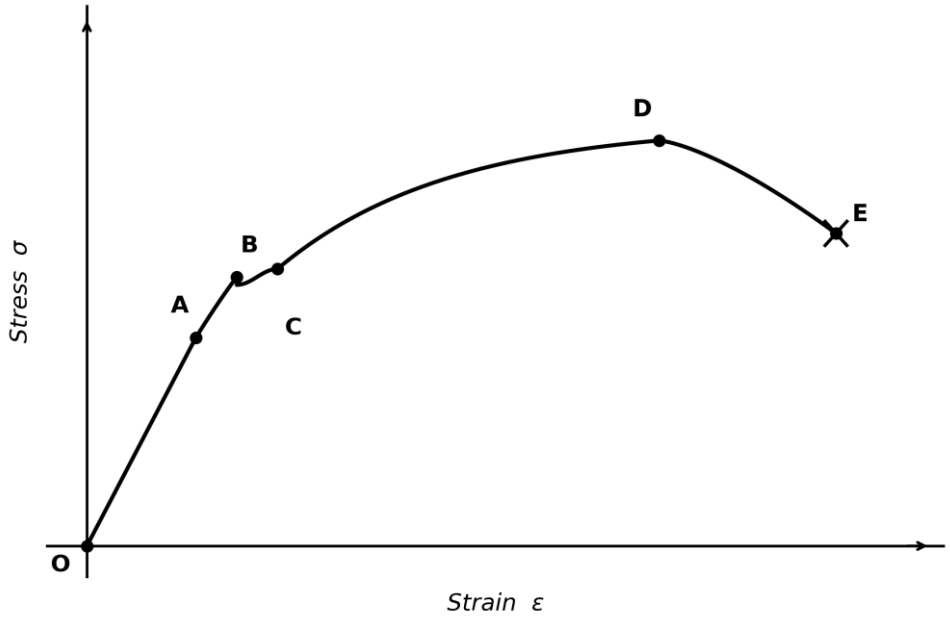


Figure: ***The stress-strain curve for a typical ductile material such as mild steel.*** *The curve rises linearly from the origin O to the proportional limit A, curves slightly to the elastic limit B, then passes the yield point C where plastic deformation begins. From C to D the material work-hardens, reaching its ultimate tensile strength at D. Beyond D the specimen necks down and the stress falls until fracture finally occurs at E, shown by the small cross.*

At the origin **O**, the material is unloaded. As we start to pull, the curve rises in a perfectly straight line from **O** up to a point marked **A**. In this region the stress is directly proportional to the strain, which is just Hooke's law in the language of stress and strain. Young's modulus of the material is the slope of this straight line, and it is the same number we would get from any elastic measurement on any piece of the same metal. Point **A** is the end of this linear behaviour. It is called the **proportional limit**, because beyond it the proportionality between stress and strain begins to break down. It can be defined as *the largest stress for which the strain remains exactly proportional to the stress.* Up to this point the material obeys Hooke's law, and the stress-strain curve is a straight line whose slope is the Young's modulus of the material.

Just beyond **A**, the curve is no longer a straight line, but the material is still elastic. If we remove the load at any point between **O** and **B**, the specimen returns to its original length and no permanent change has been made. Point **B** marks the very last stress for which this full recovery is possible. It is called the **elastic limit,** which can be defined as *the largest stress the material can sustain without suffering any permanent deformation.* Below the elastic limit, removing the load returns the specimen to its original length. Above it, some permanent deformation remains even after the load is removed. The elastic limit is marked as point **B** on the curve, and it lies a little above the proportional limit **A** for most ductile materials.

Once the stress exceeds the elastic limit, something new happens inside the metal. The molecular bonds that were stretched elastically in the earlier region are now beginning to break and reform in new positions. Planes of atoms slide past one another. Point **C** on the curve marks the moment when this sliding becomes widespread enough to produce a sudden yielding: the material gives way, and the strain increases rapidly even though the stress has barely risen (and, in some metals, has actually dipped slightly). Point **C** is called the **yield point**. Beyond it, the material has entered the plastic region, where deformation is permanent.

By definition, the **yield point** of a material *is the stress at which the material suddenly begins to deform rapidly with little or no further increase in stress.* It marks the boundary between the elastic region (where deformations are recoverable) and the plastic region (where deformations become permanent).

So **Plastic deformation** can be defined as *the permanent change in shape of a material under a load that has exceeded the elastic limit.* Unlike elastic deformation, plastic deformation remains after the load is removed, because the molecular bonds have reorganised into new positions that no longer correspond to the original unloaded shape.

From the yield point **C** onwards, the curve rises again, but in a different way. The slope is much gentler than in the elastic region, and the rise is concave, gradually levelling off. In this region the metal is strengthening itself as it deforms. Each new slip between planes of atoms leaves behind defects that make further slipping harder, so more and more stress is needed to produce each additional increment of strain. This self-strengthening process is called work hardening, and it is *the reason a*

blacksmith's hammer can turn a soft iron rod into a harder tool, and the reason a paperclip becomes stiffer the more you bend it back and forth before it finally snaps.

So **work hardening** (also called **strain hardening**) is defined as *the strengthening of a ductile material caused by plastic deformation.* As the material is deformed beyond its yield point, defects accumulate in its atomic structure that resist further deformation, so the stress required to continue stretching the material rises. Work hardening is responsible for the rising portion of the stress-strain curve between **C** and **D**.

The work-hardening region ends at point **D**, which is the highest stress the specimen ever reaches. This point is called the **ultimate tensile strength** of the material. Beyond D the curve actually starts to drop, which sounds strange because the specimen has not yet broken. The drop happens because the specimen begins to narrow locally, forming a thin neck somewhere along its length. Inside the neck the cross-sectional area is smaller than anywhere else, so for the same load the local stress in the neck rises sharply even though the total force on the specimen is no longer increasing. The plot, however, is drawn using the original (undeformed) cross-sectional area of the specimen, so it shows a falling stress, which is an artefact of using the original area rather than the true current area. Physically, the wire is being pulled apart at the neck.

Concisely, **ultimate tensile strength** of a material *is the maximum stress the material can sustain before it begins to neck down and fail.* It is the highest point reached on the stress-strain curve, marked as point D. It is also called the **breaking stress** in some textbooks, though the true fracture actually happens a little beyond this point at **E**.

The neck continues to thin until the material finally fractures at point **E**. On the graph, fracture is marked with a small cross to indicate the sudden end of the curve. The specimen is now in two pieces, and the experiment is over.

So the **fracture point** of a material can be defined as *the stress at which the specimen actually breaks into two pieces*. It occurs a little beyond the ultimate tensile strength because of the necking that takes place after D.

Three material classes

The stress-strain curve we have just walked through is typical of a ductile metal such as mild steel or copper. Not all materials behave this way. Some fracture much earlier, with almost no plastic deformation at all. Others deform enormously before they fracture, with almost no truly linear region at the start. The figure below shows the stress-strain curves of three materials drawn on the same axes for comparison, so that you can see the differences at a glance.

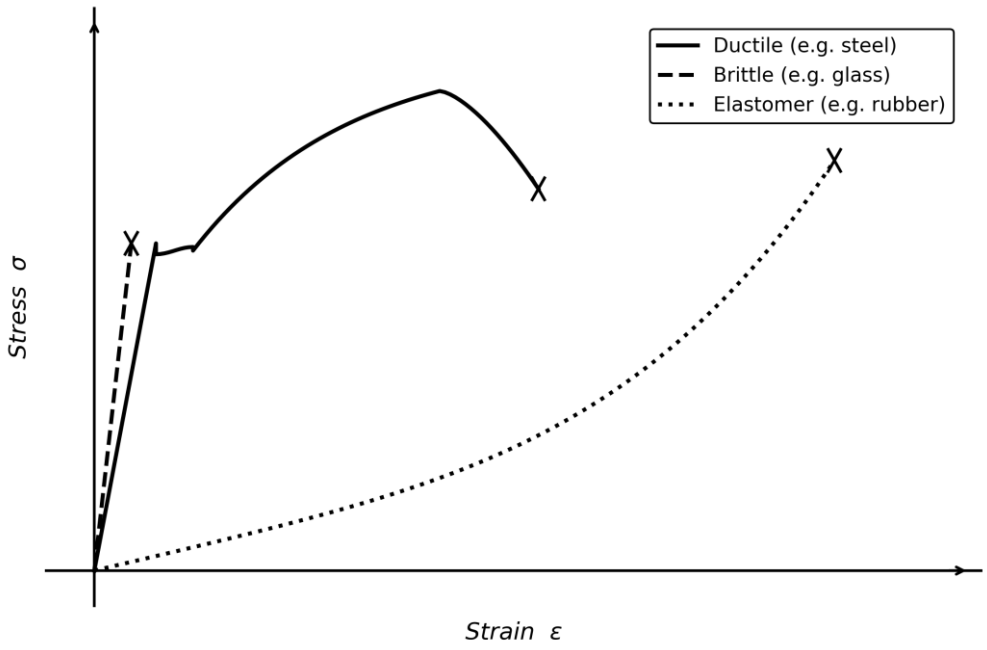


Figure: ***Stress-strain curves of three material classes drawn on the same axes for comparison.*** *The solid line shows a ductile material such as steel, with a long plastic region and eventual necking before fracture. The dashed line shows a brittle material such as glass, which fractures at a small strain with almost no plastic deformation. The dotted line shows an elastomer such as rubber, which has a very low initial slope but stiffens as the strain grows and eventually fractures at a very large strain. The small crosses indicate the fracture point of each material.*

The solid line shows the ductile curve we have already studied. Steel and copper behave this way, with a clear linear region, a clear yield point, a long plastic region with work hardening, and eventual fracture after substantial permanent deformation. Such materials are called **ductile.**

Concisely, a **ductile material** *is one that undergoes substantial plastic deformation before it fractures*. Its stress-strain curve has a long plastic region beyond the yield point, and the total strain at fracture is large compared with the strain at the elastic limit. Examples include mild steel, copper, aluminium, and most of the engineering metals.

The dashed line shows a **brittle** material such as glass or cast iron. A brittle material is stiff: its stress-strain curve rises very steeply in the elastic region. But it fractures almost immediately after the elastic limit, with almost no plastic region and almost no warning. The total strain at fracture is tiny, and the break happens suddenly. A glass rod pulled in a laboratory tensile test will snap cleanly at a small extension, with the two halves still fitting back together almost perfectly because no plastic deformation occurred to change their shapes.

Concisely, a **brittle material** *is one that fractures at or very soon after its elastic limit, with little or no plastic deformation.* Its stress-strain curve is almost entirely straight and stops abruptly at fracture. The total strain at fracture is small. Examples include glass, cast iron, ceramics, and many stones.

The dotted line shows an **elastomer** such as natural rubber. An elastomer is very different from both the ductile and the brittle classes. It has an extremely low initial slope, which means it is very soft: a small force produces a large strain. But the curve is not straight at all, and the slope steepens as the strain increases, so the material stiffens the more it is stretched. Eventually the curve becomes very steep and then fractures, often at a strain of several hundred percent, which is far beyond anything any metal could survive. A rubber band can be stretched to three or four times its natural length before it finally breaks.

Concisely, an **elastomer** *is a material that can undergo very large elastic strains (typically several hundred percent) and return to its original shape when the load is removed, but whose stress-strain relationship is strongly non-linear even at small strains.* Rubber and most soft polymers are elastomers. Elastomers do not have a proportional limit in the ordinary sense because they never have a truly linear region.

Toughness: the area under the curve

We now meet the third of the three quantities that describe the mechanical personality of a material. We already have **Young's modulus**, *which measures stiffness*, and **breaking stress,** *which measures strength*. The third quantity answers a different question altogether: *how much energy can the material absorb before it fractures?* The answer to that question is called **toughness.**

Toughness is not the same as strength, and it is not the same as stiffness. A tough material is one that absorbs a lot of energy per unit volume before giving way, regardless of whether that energy is absorbed through a small strain at high stress or through a large strain at modest stress. *On a stress-strain curve, the energy absorbed per unit volume is represented by the area under the curve from zero up to the point where fracture occurs. The tougher the material, the larger that area.*

This deserves a small pause, because we are introducing a new interpretation of the stress-strain curve. In the chapter on momentum in *Foundations of Physics* you met the idea that the area under a force-time graph equals the total impulse delivered, because impulse is defined as force multiplied by time. The same kind of reasoning applies here. The area under a stress-strain curve turns out to equal the work done per unit volume of the material, because stress is force per area and strain is extension per length, and the product of the two has the dimensions of work per volume. We shall derive this result carefully in a later part of the chapter when we study elastic potential energy, but for now the key fact is that the area under the curve is energy, and the area up to the fracture point is the total energy the material absorbed before it gave way.

Concisely, the **toughness of a material** *is the total energy per unit volume that the material can absorb before it fractures.* It is measured by the area under the stress-strain curve from zero strain up to the fracture point. Its SI unit is the joule per cubic metre (J/m^3). A tough material is not the same as a strong material: a material with a modest breaking stress but a large strain at fracture can be tougher than a material with a high breaking stress but a tiny strain at fracture.

The three material classes we just introduced have very different toughnesses, and the differences are visible directly on the combined figure. Glass, with its steep but very short curve, encloses only a very small thin (narrow) area before fracture, so its toughness is low. Steel, with its long work-hardening region stretching out to substantial strains before fracture, encloses a much larger area, so its toughness is high. Rubber, with its curve reaching out to strains of several hundred percent, encloses a truly enormous area despite its modest breaking stress, so its toughness is also very high. These three different toughnesses are exactly what Kipanga was trying to ask about in the dialogue that follows.

REAL Example 11

Mr. Akilikubwa has just finished drawing the three stress-strain curves on the board and has defined toughness as the area under the curve up to fracture. Kipanga, who has been silent for the whole demonstration, finally raises his hand.

Kipanga: *"Sir, I have been waiting for this lesson for weeks, and I think I now understand something that has been bothering me. Glass has a higher Young's modulus than rubber. Glass has a higher breaking stress than rubber. So by both stiffness and strength, glass wins. And yet when I drop a glass bottle on the floor it smashes, and when I drop a rubber ball on the floor it bounces. The glass is the winner on both scales and also the loser in the real world. How is that possible?"*

Mr. Akilikubwa: *"Kipanga, this is exactly the right question, and the answer is sitting on the board in front of you. Look at the three curves and tell me which one encloses the most area before fracture."*

Kipanga walks to the board and runs his finger along each curve.

Kipanga: *"The glass curve is very steep but stops almost immediately. The area under it is tiny. The steel curve is less steep but goes further across before it stops. Its area is bigger. The rubber curve is almost flat at the start but it stretches all the way across the board before it finally stops, so even though it is low, the total area under it is enormous."*

Mr. Akilikubwa: *"Correct. And that area is the total energy per unit volume the material can absorb before it fractures. When you drop a glass bottle on the floor, the impact delivers a certain amount of energy to the glass very quickly. That energy has to go somewhere. In a ductile or elastomeric material, it would be absorbed by the elastic and plastic deformation of the material itself, because the area under the curve provides plenty of room to absorb it. In glass, the area under the curve is so small that the incoming energy is more than the material can absorb elastically, and the only way out is to break the bonds that hold the molecules together. The crack starts at the first flaw on the surface of the bottle and races through it in microseconds, and the bottle is in pieces before you finish hearing it hit the floor."*

Mr. Akilikubwa: *"When you drop a rubber ball, the same impact delivers roughly the same amount of energy. But the rubber has a gigantic area under its curve, so the incoming energy is absorbed easily by a brief elastic deformation. The rubber compresses, stores the energy as elastic potential energy, and then releases it again by springing back to its original shape. The ball bounces, and the story ends well for the rubber."*

Kipanga: *"So the word I was missing all along was toughness, not strength."*

Mr. Akilikubwa: *"Yes. Glass is stiff and strong but not tough. Rubber is soft and weak but very tough. Steel is stiff, strong, and tough, which is why we build bridges out of it."*

Explain in your own words why a glass bottle dropped on a hard floor shatters while a rubber ball dropped on the same floor bounces back unharmed, using Young's modulus, breaking stress, and toughness as the three key quantities.

Solution

The everyday word "hard" secretly describes three independent properties, and the paradox dissolves once we give them their proper names. Young's modulus measures stiffness. Breaking stress measures strength. Toughness measures how much energy per unit volume the material can absorb before fracture. A material can rank high on any combination of these three scales, and there are materials at every corner of the three-dimensional map.

Glass is stiff (high Young's modulus) and reasonably strong (breaking stress around 50MPa) but not tough, because its stress-strain curve is almost a vertical line that stops abruptly at a very small strain. The area under the glass curve is tiny, so glass can absorb only a tiny amount of energy per unit volume before its bonds break. Dropping a bottle on the floor delivers more energy than the glass can absorb, and the only escape route is fracture.

Rubber is soft (low Young's modulus) and weaker than glass in breaking stress (around 20MPa) but is far tougher, because its stress-strain curve extends out to strains of several hundred percent before fracture. The area under the rubber curve is therefore enormous, and the rubber can absorb the energy of a normal impact without any damage at all. It deforms elastically, stores the energy briefly as elastic potential energy, and releases it again when the ball bounces back.

The everyday observation that glass breaks and rubber bounces is therefore an observation about toughness, not about stiffness or strength. Without the three separate quantities, the observation looks like a paradox. With the three separate quantities, it is an obvious consequence of where each material sits on the toughness axis.

Making Sense of the Answer: *Notice that the resolution of the paradox did not require any of the three properties to be denied. Glass really is stiffer than rubber, and glass really does have a higher breaking stress than rubber. Those two facts remain true. What was missing was the third quantity that the everyday word "hard" was silently conflating with the other two. Once toughness is added to the vocabulary, all three materials fit naturally into a consistent picture, and the puzzle of why a glass bottle breaks and a rubber ball does not stops being a puzzle at all.*

Think Like a Physicist: *Whenever everyday intuition seems to contradict a physical principle, ask whether the everyday intuition is actually using a single word to describe two or three separate physical quantities. The same trap appears with "heat" and "temperature" in thermodynamics, with "weight" and "mass" in mechanics, and with "hardness" in materials science. In each case, the resolution is not to argue about the intuition but to replace the ambiguous word with the precise physical quantities it was secretly covering.*

Elastic hysteresis and elastic fatigue

The stress-strain curves we have drawn so far all assumed that we were increasing the load from zero up to fracture. But real engineering materials are often loaded and unloaded repeatedly rather than monotonically. A spring inside a daladala bench is compressed and released thousands of times a day. A car tyre is squeezed and released with every rotation of the wheel. A rubber engine mount absorbs vibration by being cyclically loaded at high frequency. To describe these situations, we need to follow what happens when the load is removed and re-applied, and for that we need a slightly different kind of graph.

If you load a piece of rubber, then release the load slowly and watch the stress-strain curve trace itself out, you find that the curve does not come back along the same path it went up by. The unloading curve lies below the loading curve, and the two together form a closed loop as shown in the figure below. This loop is called an elastic hysteresis loop, and it tells us something important about the energy that went into stretching the rubber.

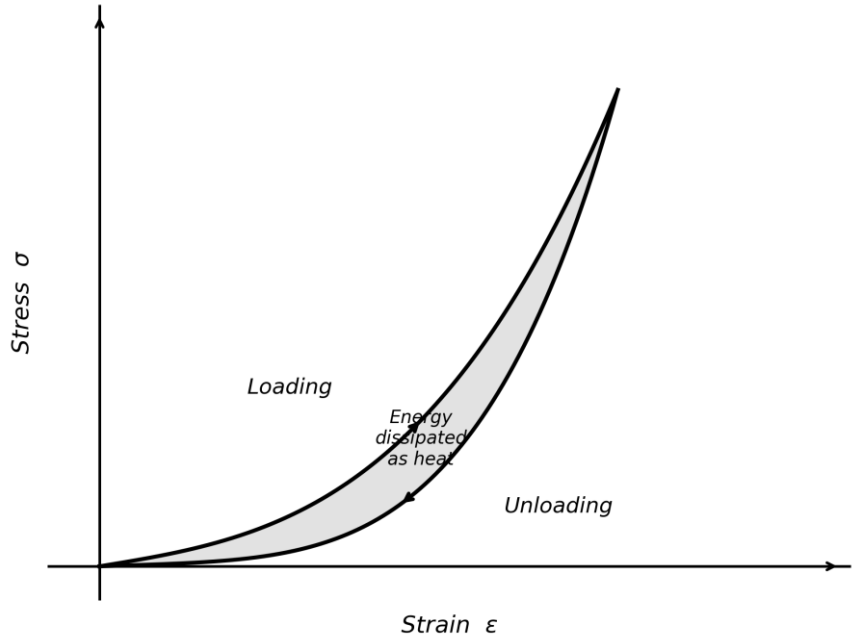


Figure: ***The elastic hysteresis loop of a rubber-like material.*** *The loading curve lies above the unloading curve, and the two together enclose a shaded area. The area under the upper curve is the energy per unit volume delivered to the rubber during stretching, and the area under the lower curve is the energy per unit volume returned when the load is removed. The shaded difference between the two areas is the energy per unit volume dissipated as heat during each complete loading and unloading cycle.*

Because the area under a stress-strain curve is energy per unit volume, the area under the upper (loading) curve is the energy per unit volume that was delivered to the rubber while it was being stretched. The area under the lower (unloading) curve is the energy per unit volume that the rubber gave back when the load was removed. These two areas are not equal: the loading area is larger. The difference between them is the shaded region enclosed by the loop, and that difference is the energy per unit volume that was lost during the cycle. The lost energy does not disappear. It is converted into heat inside the rubber, and if you load and release the rubber rapidly enough, the temperature of the rubber actually rises in a way you can feel with your hand. This is why a rubber band snapped repeatedly between your fingers becomes warm, and why the tyres of a car driven at high speed can get hot enough to damage themselves on a long journey.

Concisely, **elastic hysteresis** *is the phenomenon in which the unloading curve of a material lies below the loading curve, so that energy is lost during each loading and unloading cycle.* The lost energy appears as heat inside the material. The amount of energy lost per unit volume per cycle is equal to the area enclosed by the hysteresis loop on the stress-strain diagram. Rubber and other elastomers show strong hysteresis, while most metals show very little at loads below their elastic limit.

A related but distinct phenomenon occurs when a material is loaded repeatedly at stresses well below its breaking stress. One might expect such a material to survive forever, because each individual loading is not enough to cause damage. But real materials often fail after a large but finite number of cycles, even when the stress per cycle was always safely below the breaking stress. The failure is caused by tiny cracks that begin at microscopic flaws in the material and grow a little with each cycle, until eventually they are large enough to cause sudden fracture. This gradual weakening under cyclic loading is called **elastic fatigue**, which can be defined as *the progressive weakening of a material under repeated cycles of loading, even when each individual load is well below the breaking stress.* Elastic fatigue is responsible for many real-world engineering failures in aircraft wings, bridge cables, engine parts, and railway tracks, and is the reason such structures are routinely inspected for cracks and retired long before the number of cycles they have survived would suggest.

Elastic fatigue is also the reason a piece of wire that has been bent back and forth many times eventually snaps at the bend, even though the single bending force is much less than the force that would snap the wire if applied only once. Each bend does not break the wire, but each bend grows the tiny cracks inside the metal a little bit, until after enough cycles the cracks are large enough to cause failure. A student who wants to snap a thick wire without access to wire cutters knows this trick instinctively, even without knowing the physics: bend the wire back and forth many times, and eventually it will break.

The ideas have finished cooking; now let us sit down and enjoy them through three worked examples. The first attaches a number to the hysteresis loop just discussed, and the next two return to the stress-strain curve itself.

BINDER Example 12

A rubber engine mount is designed to absorb vibration by being cyclically compressed and released. Its stress-strain hysteresis loop encloses an area of $4 \times 10^4\ \mathrm{J/m^3}$ per cycle. The mount has a volume of $50\mathrm{cm^3}$ and is cycled 25 times per second by the running engine. Calculate (a) the energy dissipated as heat per cycle, and (b) the rate at which heat is generated inside the rubber.

Solution

(a) The energy dissipated as heat per cycle is the area enclosed by the hysteresis loop multiplied by the volume of the rubber. Calling the energy per cycle Q_1 (where Q_1 stands for the heat per cycle), the loop area u and the volume V:

$$Q_1 = u \times V$$

Substituting in SI units and computing:

$$Q_1 = (4.0 \times 10^4\ \mathrm{J/m^3})(50 \times 10^{-6}\ \mathrm{m^3}) = 2\mathrm{J}$$

Therefore, the rubber dissipates 2J of heat in each compression-and-release cycle.

(b) The rate of heat generation is the energy per cycle multiplied by the number of cycles per second. Calling the cycling frequency f and the rate of heat generation P:

$$P = Q_1 \times f$$

Substituting and computing:

$$P = (2.0\ \mathrm{J})(25\ \mathrm{s^{-1}}) = 50\mathrm{W}$$

Therefore, the rubber generates heat at a rate of 50W.

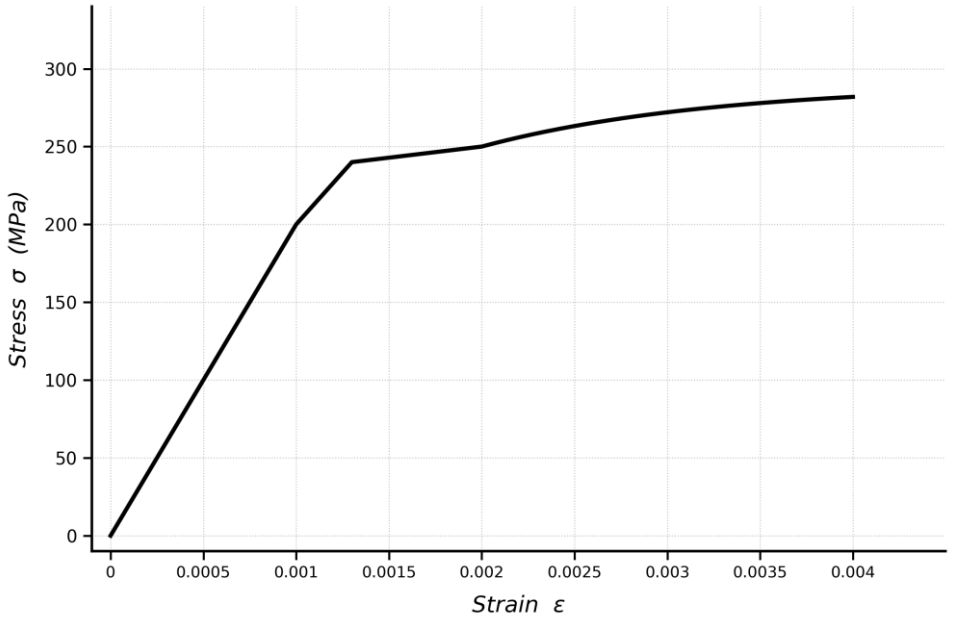
Making Sense of the Answer: *The number 50W tells us why rubber engine mounts have to be designed with cooling in mind. The heat produced per cycle is small, only 2J, but the engine delivers that small heat twenty-five times every second, and the rubber is a poor conductor that cannot shed the heat as quickly as it is being made. An undersized mount can therefore fail by overheating long before it fails by mechanical fatigue, even though it never approaches its breaking stress. Hysteresis is a quiet but persistent thief of mechanical energy, and the area inside the loop is the receipt for everything it has stolen.*

Think Like a Physicist: *Whenever a material is loaded cyclically and there is no obvious external heat source, look for hysteresis. The same calculation applies, with the same three numbers in the same three places, to brake pads on a long descent, to the rubber bushes of a daladala suspension on a Dar es Salaam pothole, and to the tyres of a long-distance lorry on a hot afternoon. In each case, the heat is made by the material from itself, and the rate is loop area times volume times frequency. Once the formula is in your toolkit, you will start to notice it everywhere mechanical work is being repeatedly done and undone.*

With a number now attached to the hysteresis loop, we return to the stress-strain curve itself and learn to read it carefully.

BINDER Example 13

The figure below shows the early portion of the stress-strain curve of a ductile metal loaded in tension. From the figure, read off (a) the proportional limit, (b) the elastic limit, (c) the stress at the yield point, and (d) estimate Young's modulus of the metal.



Solution

(a) The curve is a perfect straight line from the origin up to the point where it first begins to bend. Reading across from that point to the vertical axis gives a stress of approximately 200MPa, and reading down to the horizontal axis gives a strain of approximately 0.001.

The proportional limit of the metal is therefore about 200MPa.

(b) Just beyond the proportional limit the curve continues to rise but with a gentle curvature. The elastic limit is the highest stress for which the strain is fully recovered when the load is removed, and on the graph it sits slightly above the proportional limit, at approximately 240MPa at a strain of about 0.0013.

The elastic limit of the metal is therefore about 240MPa.

(c) The yield point is the stress at which the curve flattens noticeably, indicating that the material has begun to flow plastically. On the graph this happens at approximately 250MPa at a strain of about 0.002.

The yield point is therefore about 250MPa.

(d) Young's modulus is the slope of the initial linear portion of the curve, measured between any two points in the straight-line region. Using the origin and the proportional limit itself:

$$\mathrm{E} = \frac{\Delta\sigma}{\Delta\varepsilon} = \frac{200\mathrm{MPa} - 0}{0.001 - 0}$$

Substituting in SI units and computing:

$$\mathrm{E} = \frac{200 \times 10^{6}\mathrm{Pa}}{0.001} = 2 \times 10^{11}\mathrm{Pa} = 200\mathrm{GPa}$$

Young's modulus of the metal is 200GPa.

Making Sense of the Answer: *Every part of this question was answered by reading values directly off the graph. No extra information was needed, because the stress-strain curve contains the full elastic personality of the material in a single compact picture. This is why the curve is so valuable to engineers. Once you have it, you can find the proportional limit, the elastic limit, the yield point, Young's modulus, the ultimate tensile strength, and the breaking strain all from the same figure, without having to run any further experiments.*

Think Like a Physicist: *When you estimate Young's modulus from the slope of a graph, always use two points that are as far apart as possible within the straight-line region. A slope computed between two points that are close together is extremely sensitive to small reading errors. A slope computed between two points that are far apart averages those errors out. In this example the straight-line region happened to extend all the way from the origin to the proportional limit, so the origin and the proportional limit were the natural two points to use.*

HOT Example 14

Three unknown materials, labelled X, Y, and Z, are tested in tension in a laboratory. For each material, the load and extension are measured at several points up to the moment of fracture. From the measurements, the stress and strain are computed at each point. The tabulated results are shown below. Classify each material as ductile, brittle, or elastomeric, and justify your classification with specific reference to the shape of its stress-strain relationship.

Material X

Strain ε	Stress σ (MPa)
0.0005	100
0.001	200
0.0012	235
0.002	250
0.020	350
0.080	420
0.150	410
0.200	fracture

Material Y

Strain ε	Stress σ (MPa)
0.0002	14
0.0004	28
0.0006	42
0.0008	56
0.001	fracture

Material Z

Strain ε	Stress σ (MPa)
0.5	0.6
1.0	1.4
1.5	2.8
2.0	5.5
2.5	11.0
3.0	20.0
3.2	fracture

Solution

Material X. The first three data points lie on a straight line, because the stress rises by equal amounts for each equal step of strain: from 0 to 200MPa as the strain increases from 0 to 0.001, which is a perfect linear relationship. The fourth and fifth points show a small rise in stress for a relatively large jump in strain (235MPa to 250MPa while strain increases tenfold), which is the yield region. The remaining points show a gradual work hardening up to about 420MPa at a strain of 0.08, followed by a slight drop as the specimen necks down before fracturing at a strain of 0.20. The curve therefore has a clear linear region, a clear yield point, a long plastic region with work hardening, and substantial strain at fracture. This is the signature of a ductile material like steel.

Material Y. All four measured points lie on a straight line, because the stress rises by exactly 14MPa for each 0.0002 step of strain, all the way up to fracture at 0.001 strain. There is no yield point, no plastic region, and no work hardening. The specimen fractures abruptly at the end of the linear region, at a very small total strain. This is the signature of a brittle material like glass.

Material Z. The stress-strain relationship is non-linear from the very first measurement. A strain of 0.5 (which is already 50 percent) gives only 0.6MPa of stress, while an additional strain of 0.5 (up to 1.0, or 100 percent) gives only another 0.8MPa of stress. But from there the curve steepens sharply: each successive step of strain produces a larger increase in stress than the previous one. Fracture occurs at a strain of about 3.2, which is 320 percent of the original length. No metal could possibly survive such a strain. This is the signature of an elastomer like a natural rubber or a similar polymer.

Making Sense of the Answer: *The three materials fit the three classes cleanly, with no ambiguity. The key observations were: Material X had a clear linear region followed by a long plastic region and large strain at fracture. Material Y had only a linear region and fractured almost immediately. Material Z was non-linear from the start and stretched to several times its original length before fracturing. Each class is recognisable from the shape of the data alone, without needing to know what the material is.*

Think Like a Physicist: *When you are given tabulated laboratory data and asked to classify a material, the first thing to check is whether the first few data points lie on a straight line, which tells you whether the material has a linear elastic region at all. If it does, you can compute the slope to get Young's modulus and compare with the table of known values. Then look at how far the curve extends past the linear region before fracture: a short distance indicates brittle behaviour, a long distance indicates ductile behaviour. If there is no straight-line region at all and the strain values are of order one rather than one thousandth, the material is an elastomer.*

The stress-strain curve is now in your toolkit, along with the three quantities that describe the mechanical personality of any material: stiffness measured by Young's modulus, strength measured by breaking stress, and toughness measured by the area

under the curve. But we have so far only discussed one kind of stress and one kind of strain, namely the linear tensile kind in which force is pulled along the length of a wire. Real materials can be loaded in two other geometrically different ways: they can be sheared sideways, and they can be squeezed equally from all directions at once. Each of these other loading modes has its own stress, its own strain, and its own modulus. The next part of this chapter introduces the shear modulus and the bulk modulus, which together with Young's modulus make up the three moduli of elasticity that form the complete linear elastic description of a solid material.

SHEAR MODULUS AND BULK MODULUS

So far every stress we have met has been tensile: a force pulled on a wire along its length, and the wire stretched a little. But pulling is only one of the three geometrically distinct ways you can apply a force to a solid body. You can also push sideways on a block with a force parallel to one of its faces, and the block will skew out of shape while keeping the same volume. Or you can squeeze a body equally from all directions at once, and it will shrink in volume while keeping the same shape. Each of these three loading modes produces its own kind of stress, its own kind of strain, and its own modulus of elasticity. Tensile loading gives the Young's modulus we have already met. Sideways loading gives the **shear modulus.** All-round squeezing gives the **bulk modulus.** Between them, Young's modulus, the shear modulus, and the bulk modulus describe the complete linear elastic behaviour of any solid material.

Shear stress and shear strain

Imagine a rectangular block of material resting on a horizontal table, with its base glued firmly to the table. Now push the top face of the block with a horizontal force that is parallel to the face, not perpendicular to it. The block does not stretch or shrink in any direction, because the force is not pulling it apart or squeezing it together. Instead, the block skews out of shape. The top face slides sideways a little, the base stays where it is because it is glued, and the originally-rectangular block becomes a parallelogram. This kind of deformation is called shear, and the following figure shows what it looks like.

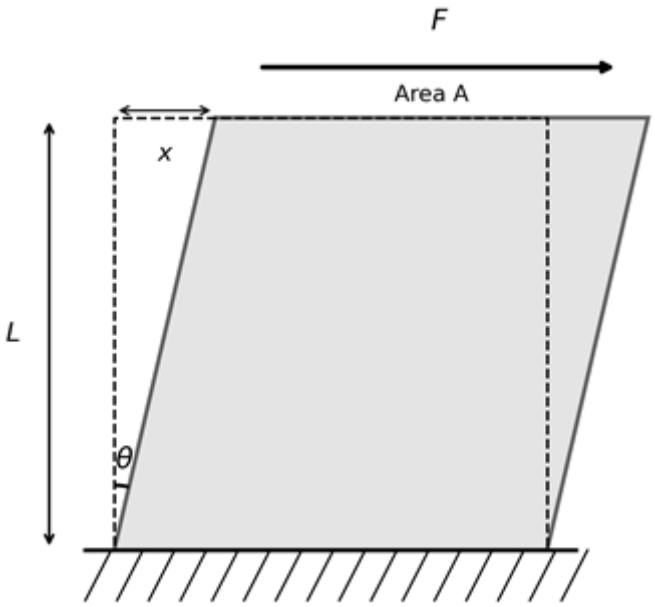


Figure: *A rectangular block of original height L with its base fixed to a rigid surface, subjected to a horizontal force F parallel to its top face. The top face, of area A, displaces sideways by a distance x while the base stays where it is. The originally rectangular block deforms into a parallelogram. The small angle θ between the original vertical side and the new slanted side is the geometric tilt that we shall use to define the shear strain in the prose below. The dashed outline shows the undeformed shape for comparison.*

To describe this deformation quantitatively we need two new quantities. The **first** is the stress produced by the **tangential force.** Just as tensile stress was defined as the normal force divided by the area of the face it acts on, **shear stress** is defined as *the ratio of the tangential force applied to a surface to the area of that surface*. Its symbol is τ (tau), and thus:

$$\boldsymbol{\tau = \frac{F}{A}}$$

Its SI unit is the pascal (Pa), the same as tensile stress. It measures how hard each unit of area is being pushed sideways, rather than pulled or squeezed.

The **second** quantity is the strain produced by the deformation. When the top face of the block slides sideways by a distance x while the base remains fixed, the originally vertical sides of the block tilt away from the vertical by a small angle. Call that angle θ, just as you would call any other angle in geometry. From the right triangle formed by the original vertical edge, the

new slanted edge, and the horizontal displacement of the top, the tangent of θ is the sideways displacement divided by the original height of the block:

$$\tan\theta = \frac{x}{L}$$

For small deformations, which is the only regime where Hooke's law applies anyway, the tangent of a small angle in radians is essentially equal to the angle itself, so we can drop the tangent and write the angle directly:

$$\theta = \frac{x}{L}$$

This small angle θ is the geometric description of the deformation, and any student of geometry could have written it down without knowing any elasticity at all. But in elasticity it has a special name and a special symbol, because it plays the same role for shear loading that the ordinary tensile strain ε plays for tensile loading. The elastic-theory name for this angle is the shear strain, and its standard symbol is **γ** (gamma). The shear strain is identical in magnitude to the tilt angle:

$$\boldsymbol{\gamma = \theta = \frac{x}{L}}$$

So the **shear strain** can be defined as *the small angle (in radians) by which the originally vertical sides of a sheared block tilt under the applied tangential force*. Equivalently, *it is the ratio of the sideways displacement of the top face of the block to the original perpendicular distance between the top face and the fixed base.* Its symbol is γ and it has no units.

The shear modulus

Now, just as tensile stress divided by tensile strain gave us Young's modulus, shear stress divided by shear strain gives us a new material property called the shear modulus. Its symbol is G and its mathematical definition is:

$$\boldsymbol{G = \frac{\tau}{\gamma}}$$

Substituting the definitions of τ and γ gives the expanded form:

$$\boldsymbol{G = \frac{\frac{F}{A}}{\frac{x}{L}} = \frac{FL}{Ax}}$$

The **shear modulus** of a material can be defined as *the ratio of the shear stress acting on the material to the resulting shear strain, provided the material is within its elastic range.* Its SI unit is the pascal (Pa). It measures how strongly the material resists being skewed out of shape. A material with a large shear modulus is hard to shear; a material with a small shear modulus is easy to shear. The shear modulus is also known as the **modulus of rigidity**, and some older textbooks use the symbol η (eta) instead of G.

The shear modulus of a typical metal is roughly one third of its Young's modulus. For steel, Young's modulus is about 200GPa and the shear modulus is about 80GPa. For copper, the corresponding values are 120GPa and 45GPa. For aluminium, 70GPa and 26GPa. The rough one-third ratio is not a coincidence: it can be derived from the theory of linear elasticity for an isotropic solid, together with one additional material parameter called Poisson's ratio, which we shall not pursue further in this chapter.

Bulk stress and bulk strain

The third and final loading mode is all-round compression. Imagine taking a small cube of material and squeezing it equally from every direction at once, the way the water pressure at the bottom of the Indian Ocean squeezes every cubic centimetre of a submarine pressure hull. The cube does not change shape, because the squeezing is equal in all directions. But it does shrink in volume, because every molecule in the interior is being pushed slightly closer to its neighbours. The following figure shows a cube subjected to such a uniform pressure.

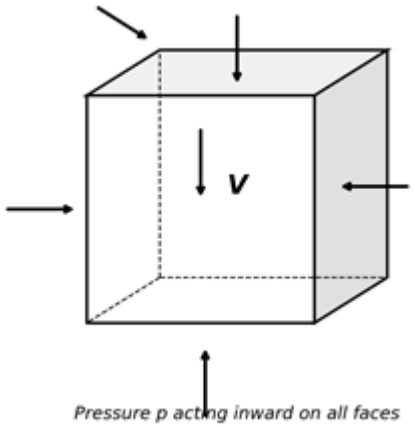


Figure: *A cube of material of volume V subjected to a uniform pressure p acting inward on all six of its faces simultaneously. The cube shrinks slightly in volume by an amount ΔV while retaining its original shape. This kind of loading is called bulk compression, and unlike tensile loading or shear loading it can be applied to gases and liquids as well as to solids.*

The stress in this situation is the pressure itself, which is the normal force per unit area acting on any face of the cube. The strain is the fractional change in volume of the cube. If the original volume of the cube is V and the volume change (a decrease, because the cube shrinks) is ΔV, then the bulk strain is:

$$\textbf{bulk strain} = -\frac{\Delta V}{V}$$

The negative implies that when the pressure on the cube increases, the volume shrinks The bulk strain is dimensionless, just like tensile strain and shear strain. The bulk strain of solids is usually a small number. A bulk strain of -0.001 means the volume has decreased by one part in a thousand under the applied pressure, and for most solids under ordinary conditions the strain is this small or smaller in magnitude.

Concisely:

Bulk stress *is the pressure acting uniformly and inward on all faces of a body*. Its symbol is p (for pressure) and its SI unit is the pascal (Pa). Unlike tensile stress and shear stress, bulk stress is the only kind of stress that can be applied to a liquid or a gas, because liquids and gases cannot sustain tensile or shear loads at rest.

Bulk strain *is the fractional change in volume of a body under a uniform pressure*. It is dimensionless, and for small deformations (the only regime where Hooke's law applies) it is always a small negative number.

The bulk modulus

The ratio of bulk stress to bulk strain is called the bulk modulus, K, and it has its own subtlety. An increase in pressure always produces a decrease in volume, so the volume change ΔV is inherently negative when the pressure change is positive. To keep the bulk modulus itself a positive quantity (which is what we want, because a positive modulus describes a stiff material and a negative modulus would describe something that expands when you squeeze it, which no ordinary material does), we include an explicit minus sign in the definition:

$$\mathbf{K} = \frac{\textbf{Bulk stress}}{\textbf{Bulk strain}} = -\frac{\Delta p}{\frac{\Delta V}{V}} = -V\frac{\Delta p}{\Delta V}$$

The SI unit of K is the pascal, the same as Young's modulus and the shear modulus.

A material with a large bulk modulus is hard to squeeze; a material with a small bulk modulus is easy to squeeze. Water has a bulk modulus of about 2.2GPa, which is why water at the bottom of the deepest ocean trench is only about 4 percent denser than water at the surface, despite the enormous pressure.

The reciprocal of the bulk modulus has its own name and symbol. It is called the **compressibility** of the material, and it measures how easily the material can be squeezed rather than how strongly it resists. Its symbol is κ (kappa) and its SI unit is the inverse pascal (1/Pa or Pa^{-1}).

$$\kappa = \frac{1}{K} = \frac{\textbf{Bulk strain}}{\textbf{Bulk stress}}$$

A highly compressible material is one whose volume changes substantially under a modest pressure. A barely compressible material is one whose volume is almost unchanged even under large pressures. Gases are highly compressible, liquids are barely compressible, and solids are almost completely incompressible by comparison.

Values of E, G, and K for common materials

The three elastic moduli of a material are related, but they are not equal to one another, and for any given material the three values are usually of the same order of magnitude while differing by factors of two or three. The table below lists typical values for the materials most commonly met in this chapter. Notice that for each metal, the shear modulus is roughly one third of Young's modulus, and the bulk modulus is comparable to Young's modulus (larger for some metals, smaller for others such as steel). For rubber, the shear modulus is tiny but the bulk modulus is comparable to water, because rubber is soft in shear but nearly incompressible.

Material	E (GPa)	G (GPa)	K (GPa)
Steel	200	80	160
Copper	120	45	140
Aluminium	70	26	76
Glass	70	28	40
Water	—	—	2.2
Rubber	0.01	0.003	2.0

Water and rubber are the two unusual entries in the table. Water has no Young's modulus and no shear modulus because it is a fluid: it cannot sustain a tensile pull or a sideways push when at rest, because it simply flows out of the way. It does, however, have a perfectly well-defined bulk modulus, because squeezing it equally from all directions cannot be avoided by flowing. Rubber, at the other end, has a shear modulus so small (three megapascals, about ten thousand times smaller than glass) that it deforms dramatically under even the gentlest sideways push, but its bulk modulus is almost as large as that of water, which makes it almost incompressible in volume. This is why a rubber ball squashed between your hands changes shape enormously but barely changes volume at all.

Why fluids have only a bulk modulus

A solid sitting on a table does not flow, because its molecules are held in fixed positions by strong intermolecular forces. When you push sideways on the top of a solid block, those forces resist the sideways push, and the block skews into a parallelogram whose shape is determined by the shear modulus. But the molecules of a fluid (either a liquid or a gas) are not held in fixed positions. They can slide past one another freely, and they will do so whenever the slightest sideways force is applied, at a rate determined by the viscosity of the fluid rather than by any elastic property. A fluid at rest therefore cannot sustain a shear stress: if you try to shear it, it flows out of the way rather than deforming elastically. The same argument applies to tensile loading: a liquid in an open container cannot be pulled apart from its surface, because the surface simply retreats from whatever is pulling on it. Only bulk compression survives, because equal inward pressure from all directions cannot be avoided by any kind of flow. This is why the table above lists a bulk modulus for water but leaves the other two columns blank, and it is why fluid mechanics needed its own chapter in ***Foundations of Physics***, with its own concepts of flow, viscosity, and pressure, separate from the elastic theory of solids that we are developing here.

For now, take a short breath and set theory aside, let us warm up with some carefully chosen worked examples.

BINDER Example 15

A rubber block of height 5cm and square top face of side 10cm has its base glued to a workbench. A horizontal force of 30N is applied to the top face parallel to one of its edges, causing the top face to displace sideways by 2mm. Calculate the shear modulus of the rubber.

Solution

The area of the top face is:

$$\mathrm{A} = (0.1\mathrm{m})(0.1\mathrm{m}) = 1 \times 10^{-2}\ \mathrm{m}^2$$

The shear stress is:

$$\tau = \frac{\mathrm{F}}{\mathrm{A}} = \frac{30\mathrm{N}}{1 \times 10^{-2}\ \mathrm{m}^2} = 3000\mathrm{Pa} = 3\mathrm{kPa}$$

The shear strain is:

$$\gamma = \frac{\mathrm{x}}{\mathrm{L}} = \frac{2 \times 10^{-3}\mathrm{m}}{5 \times 10^{-2}\mathrm{m}} = 0.04$$

The shear modulus is:

$$\mathrm{G} = \frac{\tau}{\gamma} = \frac{3000\mathrm{Pa}}{0.04} = 7.5 \times 10^{4}\mathrm{Pa} = 75\mathrm{kPa}$$

The shear modulus of the rubber block is about 75kPa, or 0.000075GPa.

Making Sense of the Answer: *Seventy-five kilopascals is an extraordinarily small shear modulus. Steel is about a million times larger at 80 gigapascals. This is why a rubber block is so easily twisted and skewed by hand, while a steel block of the same dimensions is utterly rigid to anything short of industrial machinery. Softness in shear is one of the defining features of rubber as a material, and it is exactly why rubber is used to absorb vibration in car engine mounts, machinery foundations, and earthquake-resistant buildings. The softness allows the rubber to deform freely under sideways loads, dissipating the vibrational energy harmlessly as heat instead of transmitting it to the structure above.*

Think Like a Physicist: *The shear modulus formula has exactly the same structure as the Young's modulus formula, with τ playing the role of σ and γ playing the role of ε. Any rearrangement you learned for Young's modulus works unchanged for the shear modulus, and the laboratory procedure for measuring G is essentially the same: apply a known force, measure the displacement of a known face, compute stress and strain, and divide. The only thing that changes is the geometry of how the force is applied.*

BINDER Example 16

A solid metal cube of volume 1000cm³ is lowered into the deep part of a swimming pool, where the water pressure is 1.5×10^5Pa above atmospheric. The cube shrinks in volume by 1.2mm³. Calculate the bulk modulus of the metal.

Solution

The cube shrinks, so the volume change is negative. Converting to SI units:

$$\mathrm{V} = 1000\ \mathrm{cm}^3 = 1 \times 10^{-3}\ \mathrm{m}^3$$
$$\Delta\mathrm{V} = -1.2\ \mathrm{mm}^3 = -1.2 \times 10^{-9}\ \mathrm{m}^3$$

The bulk strain is negative:

$$\frac{\Delta V}{V} = \frac{-1.2 \times 10^{-9}\text{m}^3}{1 \times 10^{-3}\text{m}^3} = -1.2 \times 10^{-6}$$

Applying the bulk modulus formula with the explicit minus sign:

$$K = -\frac{\Delta p}{\frac{\Delta V}{V}} = -\frac{1.5 \times 10^5\text{Pa}}{-1.2 \times 10^{-6}}$$

$$K = 1.25 \times 10^{11}\text{Pa} = 125\text{GPa}$$

The bulk modulus is 125GPa.

Making Sense of the Answer: *Notice how small the fractional volume change was: about one part in a million under the pressure of a few metres of water. For a metal cube of ordinary size the actual change in volume is barely measurable. This is why metals are treated as almost incompressible in most engineering calculations, and why the bulk modulus only becomes practically important in situations involving enormous pressures such as the deep ocean, high-pressure hydraulic systems, or the interior of planets.*

Think Like a Physicist: *The minus sign in the bulk modulus formula is doing real work. The volume change ΔV is genuinely negative when the pressure change is positive, so the ratio ΔV/V is also negative. Multiplying by the explicit minus sign flips the answer back to positive, which is what we want K to be for any ordinary material. In numerical problems you can track the signs carefully and the two minus signs will cancel, or you can substitute the magnitudes of ΔP and ΔV directly and get the same positive answer. Either way works, provided you understand what each minus sign is doing and do not drop any of them by accident. If you ever get a negative bulk modulus as your final answer, you have made a sign error somewhere.*

REAL Example 17

Kipute is reading the back of her uncle's diving wristwatch, which he uses on weekend dives off the coast of Mafia Island. The case is engraved with the words "water resistant to 200m". The case is made of stainless steel with a bulk modulus of 160GPa, and its internal volume is 5cm³. Take the density of seawater as 1025kg/m³ and $g = 9.8\text{m/s}^2$. Find the change in the internal volume of the watch case at a depth of 200m.

Solution

The gauge pressure at a depth of 200m below the surface is:

$$\Delta p = \rho g h = (1025\text{kg/m}^3)(9.8\text{m/s}^2)(200\text{m}) = 2.01 \times 10^6\text{Pa}$$

Rearranging the bulk modulus formula for the bulk strain gives $\frac{\Delta V}{V} = -\frac{\Delta p}{K}$, where the minus sign tells us the volume decreases when the pressure rises. Computing:

$$\frac{\Delta V}{V} = -\frac{2.01 \times 10^6\text{Pa}}{160 \times 10^9\text{Pa}} = -1.26 \times 10^{-5}$$

The actual change in the internal volume of the case is:

$$\Delta V = V \times \frac{\Delta V}{V} = (-1.26 \times 10^{-5})(5\text{cm}^3) = -6.3 \times 10^{-5}\ \text{cm}^3$$

The negative sign confirms that the case shrinks, and the magnitude of the shrinkage is 6.3×10^{-5} cm³ at a depth of 200m.

Making Sense of the Answer: *Two hundred metres of seawater is a serious depth. The pressure at that depth is twenty times atmospheric, and the diver feels it on every square centimetre of skin. Yet the internal volume of the steel watch case shrinks by only six hundredths of a millionth of a litre, which is so small that no instrument the diver carries could measure it directly. Steel is so stiff against bulk compression that the case is, for all practical purposes, rigid. The gut feeling that the case might be crushed by the pressure is wrong by many orders of magnitude. What actually limits the depth rating of a diving watch is not the compressibility of the steel case but the seal between the case and the glass face, which is what the manufacturer is really certifying when they engrave "water resistant to 200m" on the back.*

Think Like a Physicist: *When you meet a problem about something being crushed, broken, or deformed by water pressure at depth, always compute the actual fractional change before you trust your intuition. Water pressure feels enormous when you experience it on your eardrums, because biological tissue is much more compressible than metal. But for a metal object of any reasonable size, even the pressure at the deepest part of the Indian Ocean produces a fractional volume change of less than a tenth of a percent, which is far smaller than the precision of most engineering tolerances. The deepest point of the* ***Mariana Trench*** *(the deepest part of the world's oceans, found in Pacific Ocean), where the pressure exceeds a thousand atmospheres, would shrink a steel cube by only about half a percent of its volume.*

HOT Example 18

A copper cube of side 5cm is subjected simultaneously to two experiments in the same laboratory. In the first experiment, its base is glued to a bench and a horizontal force of 2000N is applied to its top face, parallel to one edge. In the second experiment, it is placed in a high-pressure chamber and subjected to a uniform pressure of 1×10^8Pa above atmospheric. Take the shear modulus of copper as 45GPa and the bulk modulus of copper as 140GPa. Find (a) the sideways displacement

of the top face in the shear experiment, and (b) the fractional change in volume of the cube in the bulk experiment. Comment on why the two answers are so different.

Solution
(a) Shear experiment.
The area of the top face is:

$$A = (0.05\text{m})(0.05\text{m}) = 2.5 \times 10^{-3}\,\text{m}^2$$

The shear stress is:

$$\tau = \frac{F}{A} = \frac{2000\text{N}}{2.5 \times 10^{-3}\,\text{m}^2} = 8 \times 10^5\text{Pa}$$

The shear strain is:

$$\gamma = \frac{\tau}{G} = \frac{8 \times 10^5\text{Pa}}{45 \times 10^9\text{Pa}} = 1.78 \times 10^{-5}$$

The sideways displacement of the top face is:

$$x = \gamma L = (1.78 \times 10^{-5})(0.05\text{m}) = 8.9 \times 10^{-7}\text{m}$$

The top face moves sideways by about 0.89 micrometres, which is less than a thousandth of a millimetre.
(b) Bulk experiment.
The fractional change in volume is:

$$\frac{\Delta V}{V} = -\frac{p}{K} = -\frac{1 \times 10^8\text{Pa}}{140 \times 10^9\text{Pa}} = -7.1 \times 10^{-4}$$

The negative sign confirms that the cube shrinks. In magnitude, the volume decreases by about 0.07 percent, which is a far larger fractional deformation than the shear strain of part (a).

Comment. The two experiments were carried out on the same cube of the same metal, yet the fractional deformation in the bulk experiment was about forty times $\left(\frac{7.1 \times 10^{-4}}{1.78 \times 10^{-5}}\right)$ larger than the strain in the shear experiment. The reason is that the pressure applied in the bulk experiment ($1 \times 10^8\text{Pa}$, or 100MPa) was much larger than the shear stress in the shear experiment (0.8MPa). A small force applied tangentially produces a modest shear stress; an enormous pressure applied equally from all directions produces an enormous bulk stress.

Making Sense of the Answer: *The two loading modes are of comparable physical importance for copper, but they show up in different circumstances. Shear loading matters when a metal part is being twisted, bolted, or punched, which is almost every engineering situation involving metal. Bulk loading matters only when the surrounding pressure is extreme: deep ocean, hydraulic cylinders, the interiors of planets, or the compression chambers of diesel engines. For everyday structural calculations, the shear modulus is the one you will meet far more often.*

Think Like a Physicist: *Whenever a problem involves two independent loading modes on the same object, solve each mode as if the other did not exist. The linearity of elastic theory guarantees that the two responses do not interfere with each other, so the shear displacement in this example is exactly the same whether the cube is under pressure or not, and the volume change is exactly the same whether the cube is being sheared or not. This independence is the principle of superposition, and it holds for any linear response theory in physics, including electromagnetism, small-oscillation mechanics, and low-amplitude wave motion.*

Three moduli, three loading modes, one complete description of how a solid responds to small forces. *What happens to the energy delivered to the material while it is being loaded? Where does it go, and how much of it comes back when the load is removed?* The answer is in the next part of this chapter.

ELASTIC POTENTIAL ENERGY

Pull on a rubber band, and you can feel the rubber band pulling back. Let it go, and it snaps to its original length with enough energy to sting your fingers if they are in the wrong place. The energy that was delivered to the rubber band while you were stretching it did not vanish into the air and it did not turn into heat. It was stored inside the rubber, waiting. The moment you let go, the rubber returned every last joule of it.

This is the question we closed the previous part of the chapter with: *where does the energy delivered to a stretched material go while the material is stretched, and how much of it comes back when the load is removed?* The answer is that the energy is stored inside the material as elastic potential energy, and for any material within its elastic region all of it comes back. Nothing is lost, nothing is wasted, and the formula for how much is stored turns out to be one of the cleanest results in the whole of mechanics.

Deriving the energy formula

Take a spring of spring constant k and stretch it slowly from its natural length up to a final extension x. By "slowly" we mean unhurriedly enough that nothing accelerates and no kinetic energy is involved, so that every joule of work done by our hand

goes straight into the spring and stays there. At an intermediate extension x' the applied force is $F = kx'$ by Hooke's law, and extending the spring by a further dx' does work $dW = F\,dx' = kx'\,dx'$. The total work from zero extension to the final x is:

$$W = \int_0^x kx'\,dx' = k\left[\frac{x'^2}{2}\right]_0^x = \frac{1}{2}kx^2$$

Because the process was slow, no kinetic energy appeared along the way, and every joule of this work is now sitting inside the spring as **elastic potential energy U**. That gives us our first and most important result:

$$\mathbf{U = \frac{1}{2}kx^2}$$

There is a geometric way to see the same result, and it is worth a moment because it will reappear in other corners of physics. Plot the applied force against the extension. Because the spring obeys Hooke's law, the plot is a straight line through the origin with slope k, as shown in the figure below. The work done in stretching the spring from zero to the final extension x is the area under this line (because work is force times distance, and the area of any thin vertical strip of the graph is exactly the small work $F\,dx'$ done over that strip of extension). The line is straight, so the area is a simple triangle of base x and height $kx\left(\frac{y}{x} = k\right)$. Its area is:

$$\frac{1}{2}(x)(kx) = \mathbf{\frac{1}{2}kx^2}$$

The same result we obtained from the integral.

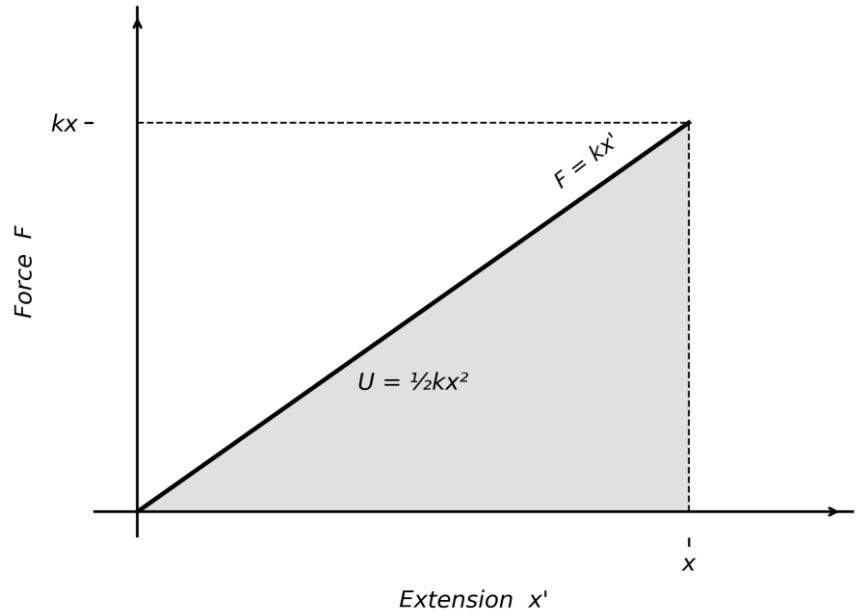


Figure: *The force required to hold a Hookean spring at extension x' is F = kx', plotted as a straight line through the origin. The work done in stretching the spring from zero to a final extension x is the area under the line, which is the shaded triangle of base x and height kx. Its area (1/2)kx² is the same result obtained by integration and represents the elastic potential energy stored in the spring.*

That triangle is not unique to springs. Any linear restoring system, when slowly loaded from rest, produces exactly the same picture: a straight-line graph, a triangular area, and a factor of one half in the stored energy formula. You will meet the same triangle again when we come to the energy stored in a charged capacitor, and again in the energy stored in a current-carrying inductor, and again in any other linear-response system where a driving quantity pulls gently against a restoring quantity. Once you recognise the triangle, half the energy problems in physics are already half solved.

Now, we can define the elastic potential energy stored in a body that obeys Hooke's law as *the work done against the restoring force in stretching it from its natural length to its present extension.*

Three equivalent forms

The formula $U = \frac{1}{2}kx^2$ is the most common form, but a problem does not always hand you the extension as the thing you know. Sometimes you know the final force instead, and sometimes you know both. Substituting $F = kx$ into the energy formula in two different ways gives two further useful forms:

$$\mathbf{U = \frac{1}{2}kx^2 = \frac{1}{2}Fx = \frac{F^2}{2k}}$$

Use the first when the extension is known, the second when both force and extension are known, and the third when only the force is known.

The final form is derived from $U = \frac{1}{2}kx^2$.

Multiplying by $\frac{k}{k}$:

$$U = \frac{1}{2}kx^2 \times \left(\frac{k}{k}\right) = \frac{(kx)^2}{2k} = \frac{F^2}{2k}$$

Stretched wire and energy per unit volume

Everything we have done so far applies to a spring, and you might worry that a wire is a different beast. It is not. Earlier in this chapter we showed that a wire of length L, cross-sectional area A, and Young's modulus E behaves as a spring of constant $k = \frac{EA}{L}$, so the spring formula carries over word for word once the spring constant is rewritten in the language of the wire. Substituting gives the energy stored in the wire when it is stretched by ΔL:

$$U = \frac{1}{2}\left(\frac{EA}{L}\right)(\Delta L)^2$$

This tells us how much energy a whole wire stores, but it depends on the dimensions of that particular wire. A more useful quantity is the energy stored per unit volume of the material, which is a pure property of the material itself and does not care about whether your sample is a short thick rod or a long thin filament. Dividing by the volume $V = AL$ and tidying up gives:

$$\frac{\mathbf{U}}{\mathbf{V}} = \frac{1}{2}\left(\frac{EA}{L \times AL}\right)(\Delta L)^2 = \frac{1}{2} \times E \times \left(\frac{\Delta L}{L}\right)^2 = \mathbf{\frac{1}{2}E\varepsilon^2} = \mathbf{\frac{1}{2}\sigma\varepsilon} = \mathbf{\frac{\sigma^2}{2E}}$$

Look at the form $\frac{U}{V} = \frac{1}{2}\sigma\varepsilon$. That is the area of a triangle with base ε and height σ, which is exactly the area under a linear stress-strain curve from zero strain up to the current strain. Earlier, when we first met the idea of toughness, we claimed that the area under a stress-strain curve represents the energy stored per unit volume, and we promised a proper derivation later in the chapter. This is that promise kept. Toughness is just this same quantity evaluated at the fracture strain, and the area-equals-energy statement that we once had to take on trust has now earned its place honestly.

Concisely, **strain energy per unit volume** of an elastically stretched material *is the elastic potential energy stored in each unit of volume*. For a material within its linear elastic region it equals $\frac{1}{2}\sigma\varepsilon$, where σ is the current stress and ε is the current strain. Its SI unit is the joule per cubic metre (J/m^3).

One formula, three forms, and a family of problems that all yield to the same triangle. Some come dressed as springs, some as wires, and some as catapults, and one or two will pretend the elastic energy has nothing to do with kinetic energy. Do not be fooled. The triangle is always the same triangle. Pick up your pen for worked examples.

BINDER Example 19

A spring of spring constant 200N/m is stretched from its natural length by 0.15m. Calculate the elastic potential energy stored in the spring.

Solution

$$U = \frac{1}{2}kx^2 = \frac{1}{2}(200\text{N/m})(0.15\text{m})^2 = 2.25\text{J}$$

The spring stores 2.25J at this extension.

Making Sense of the Answer: *Two and a quarter joules is roughly the kinetic energy of a cricket ball travelling at 7 metres per second, or the work done lifting a half-kilogram mass to waist height. The spring holds exactly that much and releases all of it the moment it is allowed to contract.*

Think Like a Physicist: *The x in U = (1/2)kx² is always the extension from the natural length, never from some other starting point such as the equilibrium position of a hanging mass. A spring hanging from a ceiling with a mass on the end has a natural length and a stretched length, and x is the difference between them.*

REAL Example 20

Kipanga has built a catapult by tying a rubber band between the two prongs of a Y-shaped stick. The rubber band behaves approximately as a Hookean spring of spring constant 80N/m. He loads the catapult with a small stone of mass 20g, pulls the rubber band back by 0.12m, and releases it. Find the speed of the stone as it leaves the catapult.

Solution

Elastic potential energy stored in the rubber band at the moment of release:

$$U = \frac{1}{2}kx^2$$

All of this converts to kinetic energy of the stone:

$$\frac{1}{2}mv^2 = \frac{1}{2}kx^2$$

$$v = \sqrt{\frac{kx^2}{m}} = x\sqrt{\frac{k}{m}} = 0.12\text{m}\sqrt{\frac{80\text{N/m}}{0.020\text{kg}}} = 7.59\text{m/s}$$

The stone leaves the catapult at approximately 7.6m/s.

Making Sense of the Answer: *Seven and a half metres per second is a realistic catapult speed, fast enough to cause a noticeable impact but slow enough to dodge if you see it coming. The entire energy budget of the shot is less than one joule, converted cleanly from elastic to kinetic by the same conservation-of-energy principle that governs every collision and every projectile.*

Think Like a Physicist: *A physicist does not compute intermediate numerical values unless the problem specifically asks for them. Writing the energy balance ½kx² = ½mv² and cancelling the halves before substituting numbers is always cleaner than computing U first and then plugging it into the kinetic energy formula, because it keeps the final answer in terms of the quantities the problem actually gave you. Symbolic first, numerical last, is the default approach. It also makes the physics visible. The formula v = x√(k/m) immediately tells you that doubling the extension doubles the speed, that a stiffer rubber band gives a faster shot, and that a heavier stone goes slower, without any arithmetic at all. Numbers hide that kind of insight behind a wall of digits; symbols put it in plain view.*

HOT Example 21

A steel wire of length 2m, cross-sectional area $1 \times 10^{-6}\ \text{m}^2$, and Young's modulus 200GPa is stretched by 1mm. Using both the spring-analogue method and the energy-density method, calculate the elastic potential energy stored in the wire and confirm that the two methods agree.

Solution

By using the spring-analogue method:

The wire is a spring in disguise, with spring constant $k = \frac{EA}{L}$, so the cleanest route to the energy is the spring formula $U = \frac{1}{2}kx^2$ with $x = \Delta L$. Substituting the wire's spring constant directly:

$$U = \frac{1}{2}\frac{EA}{L}(\Delta L)^2$$

Substituting the numbers in SI units:

$$U = \frac{1}{2}\frac{(200 \times 10^9\text{Pa})(1 \times 10^{-6}\ \text{m}^2)}{2\text{m}}(1 \times 10^{-3}\text{m})^2$$

$$U = 0.05\text{J} = 50\text{mJ}$$

The wire stores 0.05J of elastic potential energy at this extension.

By using the energy-density method:

This is done by using the energy-per-unit-volume formula $\frac{U}{V} = \frac{1}{2}\sigma\varepsilon$.

Computing the stress and strain from the given loading:

$$\sigma = E\varepsilon = E\frac{\Delta L}{L} = (200 \times 10^9\text{Pa}) \times \frac{1 \times 10^{-3}\text{m}}{2\text{m}} = 1 \times 10^8\text{Pa}$$

and the strain is $\varepsilon = \frac{\Delta L}{L} = \frac{1 \times 10^{-3}\text{m}}{2\text{m}} = 5 \times 10^{-4}$

So the energy per unit volume is:

$$\frac{U}{V} = \frac{1}{2}\sigma\varepsilon = \frac{1}{2}(1 \times 10^8\text{Pa})(5 \times 10^{-4}) = 2.5 \times 10^4\ \text{J/m}^3$$

Multiplying by the wire's volume $V = AL = (1 \times 10^{-6}\ \text{m}^2)(2\text{m}) = 2 \times 10^{-6}\ \text{m}^3$:

$$U = \frac{U}{V} \times V = (2.5 \times 10^4\ \text{J/m}^3)(2 \times 10^{-6}\ \text{m}^3) = 0.05\text{J}$$

Hence the two routes agree, giving 0.05J of elastic potential energy.

Making Sense of the Answer: $0.05J$ *packed into two cubic centimetres of wire works out to twenty five thousand joules per cubic metre. A car tyre at normal pressure stores about a kilojoule of elastic energy, which is why a ruptured tyre can injure anyone standing nearby. Elastic energy density is one reason pressure vessel design is a serious discipline.*

Think Like a Physicist: *When a calculation can be done by two different formulas that ought to agree, doing both is not double work. It is the cheapest check available, and it catches errors that a single calculation would leave hidden. More deeply, the agreement of the two methods reinforces the link between the object-level description (the total energy in this particular wire) and the material-level description (the energy per unit volume of steel, multiplied by the wire's volume). The first depends on what you built; the second depends only on what you built it from. This intensive-versus-extensive distinction appears in every later chapter of this book, and recognising it early pays dividends long after you forget the details of any particular wire.*

So far, we have three moduli, four strain formulas, and one energy equation written in six different ways. What we do not yet have is a laboratory apparatus that lets a Form Five student actually measure Young's modulus for themselves. The good news is that the apparatus is waiting for us in the next part of this chapter.

EXPERIMENTAL DETERMINATION OF YOUNG'S MODULUS

So far, Young's modulus has been a number that appears in formulas. We have used it to predict how much a wire will stretch under a given load, to compute the spring constant of a wire from its dimensions, and to calculate the elastic potential energy stored in a loaded rod. But every number in physics that deserves to be called a material property has to be measurable, and a student who has used Young's modulus repeatedly without ever having determined it for a real wire has been handed a gift without being shown the price tag. The purpose of this section is to show how that price is paid in the laboratory.

The measurement challenge

At first sight the measurement looks straightforward. Hang a known weight from a wire of known length and cross-section, measure how much the wire extends, and rearrange the equation for Young's modulus to obtain it in one step. The difficulty is hidden in the extension. For a steel wire two metres long and one millimetre in diameter, a load of fifty newtons produces an extension of less than a millimetre, often only a few tenths of a millimetre. Measuring such a small length with a ruler is impossible. Measuring it with a vernier caliper is only just possible. Measuring it accurately enough for a Young's modulus calculation requires an apparatus that was designed with this problem in mind.

The apparatus in question is Searle's two-wire arrangement. Its essential trick is that the quantity we actually need to measure is not the length of the wire itself but the change in the length of the wire. If we could somehow measure the change directly without first measuring the full length and then subtracting the initial length, most of the error would disappear. Searle's method does exactly that, and the rest of this section is devoted to understanding how it works and why every one of its features is there for a reason.

Searle's two-wire apparatus

Two identical wires are suspended side by side from the same rigid support overhead. Both wires are made from the same material and are cut from the same spool, so that they are as alike as a pair of wires can be. One of them is the test wire, which will carry the load whose effect we want to measure. The other is the reference wire, which is kept under a constant small load just heavy enough to keep it straight. A light rigid frame is attached to the lower end of each wire, and the two frames are connected to each other by a sensitive measuring device, traditionally a spirit level fitted with a micrometer screw gauge, though modern laboratory versions use a dial gauge or a digital displacement sensor in place of the spirit level.

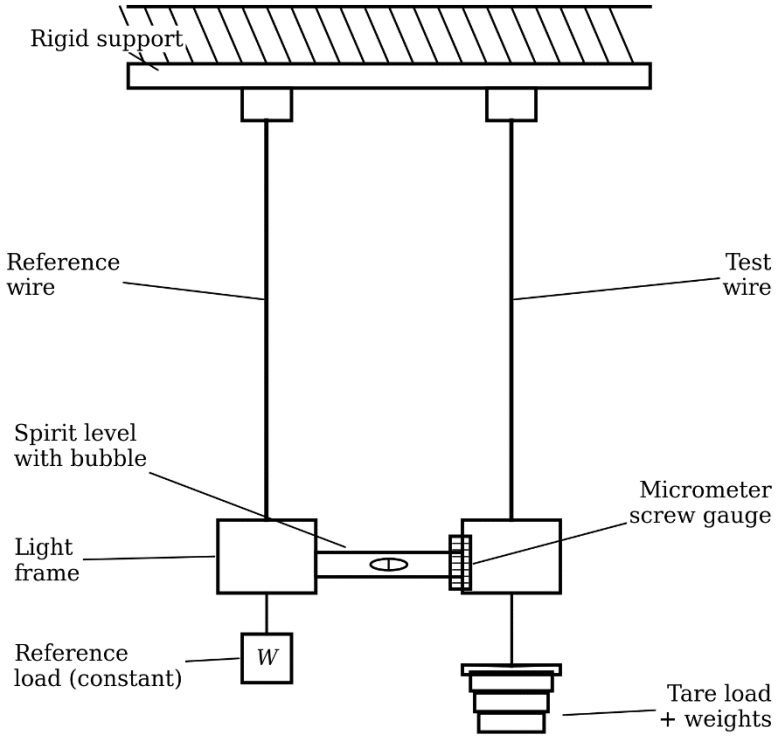


Figure: ***Searle's two-wire apparatus for measuring Young's modulus of a wire.*** *Two identical wires hang side by side from the same rigid support. The left wire is the reference wire, kept under a constant small load W just heavy enough to keep it straight. The right wire is the test wire, carrying the tare load and any additional weights whose effect on its extension is to be measured. A light rigid frame is attached to the lower end of each wire, and the two frames are connected by a sensitive*

spirit level fitted with a micrometer screw gauge. Because both wires hang in the same air at the same temperature, any thermal expansion affects them equally and cancels out of the difference in their lengths. The spirit level therefore records only the extra extension of the test wire caused by the load, which is exactly the quantity needed for the calculation of Young's modulus.

The key insight behind the two-wire arrangement is thermal compensation. Whatever expansion or contraction happens to the reference wire during the experiment happens equally to the test wire, because the two are hanging in the same air, at the same temperature, at the same moment. If the room warms up by half a degree during the experiment, both wires lengthen by the same small amount, and the spirit level is not disturbed at all. The only thing that can make the two wires differ in length is the load applied to the test wire, and that is exactly what we want to measure. Thermal effects, which would otherwise contaminate the measurement of a single wire, cancel exactly.

The reference wire plays a second, quieter role as well. It absorbs any small unintended movements of the support beam from which both wires are suspended, for example from vibration or from the slow settling of the clamp during a long experiment. Whatever happens to the beam happens to the tops of both wires equally, and therefore does not appear in the difference.

Taking the readings

The procedure begins with a small constant load, called the tare load, already hanging from the test wire. The tare load is there to remove any initial slack or kinks in the wire. Without it, the first few increments of load would go into straightening the wire rather than stretching it, and the early data would sag below the straight line we expect. With the tare load in place, the spirit level is adjusted until the bubble sits in the middle of its scale, and the initial reading on the micrometer is recorded as the zero of the experiment.

Known weights are then added to the test wire in equal increments, typically half a kilogram or one kilogram per step. After each weight is added, the spirit level tips slightly because the test wire has stretched more than the reference wire. The micrometer is turned until the bubble is centred again, and the new reading is recorded. The difference between the new reading and the initial reading is the extension produced by the total load so far. Five or six such increments give a table of load and extension values, and the experiment is concluded by carefully removing the weights and checking that the wire has returned to its original reading. If it has, the entire experiment was conducted within the elastic limit, and the data are trustworthy. If it has not, the wire was loaded beyond the elastic limit at some point, and the measurements must be repeated with smaller weights on a fresh wire.

From gradient to Young's modulus

The data are plotted with the load F on the vertical axis and the extension on the horizontal axis. Within the elastic limit, the graph is a straight line through the origin, because Hooke's law holds for the wire as a whole: the load is proportional to the extension. The gradient of the line is the force required to produce unit extension, which is exactly what we called the spring constant of the wire in the section on Young's modulus, and which we showed to equal:

$$\textbf{Gradient} = \frac{\mathbf{EA}}{\mathbf{L}}$$

Rearranging this relation for Young's modulus gives the result we are after:

$$\mathbf{E} = \frac{\mathbf{Gradient} \times \mathbf{L}}{\mathbf{A}}$$

The power of this method lies in the fact that Young's modulus is extracted from the slope of a best-fit straight line through several data points, not from a single load-extension measurement. A single measurement could be wildly wrong because of a sticky spirit-level bubble, a slipped weight, or a misread micrometer. A best-fit line through six points averages out such random errors and gives a value of Young's modulus that is far more reliable than any single-point calculation. This is why the experiment is always done as a graph, not as a single division.

Precautions and dominant sources of uncertainty

Four precautions matter more than all the others put together.

1) The initial tare load must always be present before the first reading is taken, so that the wire is being stretched rather than straightened when the extra weights begin to arrive.
2) The diameter of the wire must be measured in several places and in two perpendicular orientations at each place, and an average taken, because real wires are never perfectly uniform and never perfectly circular.
3) The experiment must never exceed the elastic limit of the wire, because once plastic deformation has begun, the wire will not return to its original length and all subsequent readings will be contaminated.
4) Parallax must be avoided when reading the spirit level, by looking at the bubble from directly above.

One of these precautions carries far more weight than the others, and it is not obvious which one until the error analysis is carried out. The reason is that the cross-sectional area of the wire depends on the square of the diameter:

$$A = \pi\left(\frac{d}{2}\right)^2 = \frac{\pi d^2}{4}$$

A small fractional error in the diameter therefore produces twice as large a fractional error in the area, and hence twice as large a fractional error in Young's modulus. When the individual uncertainties from the length, extension, and diameter are combined, the result is:

$$\frac{\Delta E}{E} = \frac{\Delta L}{L} + \frac{\Delta(\Delta L)}{\Delta L} + 2 \times \frac{\Delta d}{d}$$

Notice the factor of two in front of the diameter term. Even when the diameter is measured more precisely than any other quantity in the experiment, its contribution to the total uncertainty in Young's modulus is doubled before it is added to the sum. This is why students doing Searle's experiment are told, often with some irritation, to measure the diameter with a micrometer screw gauge and not with a vernier caliper. The reduced uncertainty in the diameter is worth more than an equivalent reduction in any of the other quantities, because the diameter enters the formula squared.

A note on the shear modulus

The experiment we have just described gives Young's modulus of a wire by a static method: weight is added, the extension is measured, the graph is plotted, and the slope is converted to Young's modulus. The corresponding experiment for the shear modulus is not static but dynamic. It uses a wire suspended vertically from a rigid support with a horizontal disc attached at the bottom, and it measures the period of torsional oscillations of the disc as the wire twists back and forth. The shear modulus is then extracted from the period. This is a simple harmonic motion experiment, and it belongs naturally to the chapter on simple harmonic motion later in this book. We shall meet it again there, once simple harmonic motion has given us the language in which the period of a torsional oscillation can be properly written.

Three worked examples now put the apparatus and the error analysis to use.

BINDER Example 22

A Form Five student performs Searle's experiment on a steel wire of length 2.00m and diameter 0.80mm. After adding a tare load and zeroing the spirit level, the student records the following five (load, extension) values:

Load F (N)	10	20	30	40	50
Extension ΔL (mm)	0.20	0.38	0.60	0.78	1.00

Plot a suitable graph, determine its gradient, and use the gradient to calculate Young's modulus of the steel wire. Compare your value with the standard value of 2.0×10^{11} Pa for steel.

Solution

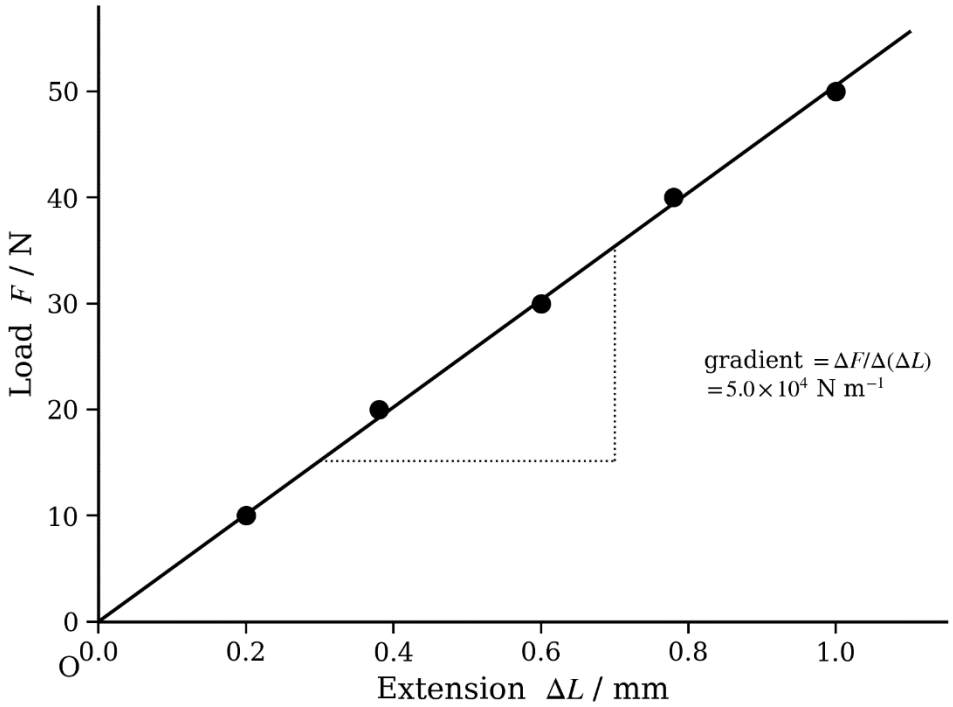


Figure: *Student's plot of load F against extension ΔL from Searle's experiment on a steel wire of length 2.00m and diameter 0.80mm.*

$$\text{Gradient} = \frac{(50 - 0)\text{N}}{(1.00 - 0)} \times 10^{-3}\,\text{m} = 5.0 \times 10^{4}\,\text{N/m}$$

The cross-sectional area of the wire is calculated from its diameter using the formula for the area of a circle.

$$A = \frac{\pi d^2}{4} = \frac{\pi \times (0.80 \times 10^{-3}\ \text{m})^2}{4} = 5.03 \times 10^{-7}\ \text{m}^2$$

Young's modulus is then obtained by multiplying the gradient by the length of the wire and dividing by the cross-sectional area:

$$E = \frac{\text{gradient} \times L}{A} = \frac{5.0 \times 10^4\ \text{N/m} \times 2.00\text{m}}{5.03 \times 10^{-7}\ \text{m}^2} = 1.99 \times 10^{11}\ \text{Pa}$$

Therefore, Young's modulus of the steel wire is approximately 2.0×10^{11} Pa, which agrees with the standard value for steel to within the precision of the experiment.

Making Sense of the Answer: *The agreement between the experimental value and the standard value is the strongest possible evidence that the apparatus, the procedure, and the arithmetic are all working correctly. Notice that the final answer comes from the slope of the straight line, not from any single data point, so a bad reading at one or two increments would not have wrecked the result. This is the quiet magic of doing physics with graphs rather than with individual measurements.*

Think Like a Physicist: *Whenever a quantity can be extracted from the slope of a straight line fitted to several data points, prefer the slope method over a single-point calculation even when the single-point answer would seem easier to compute. The slope method automatically averages out random errors in individual readings and uses every data point as evidence, whereas the single-point method stakes the entire answer on one lucky measurement. This is one of the most transferable habits in all of experimental physics, and you will use it for every experiment in this book from now on.*

REAL Example 23

In Searle's experiment for measuring Young's modulus, a small constant weight known as the tare load is always hung from the test wire before the first reading is taken. A careful student might object that this practice is unfair, because the wire is already deformed when the so-called zero of the measurement is recorded on the micrometer. Explain the purpose of the tare load, describe what the load-extension data would look like if the tare load were omitted and the experiment were started from a bare wire, and state why the extension caused by the tare load itself does not need to be known or recorded for the final calculation of Young's modulus to be correct.

Solution

The purpose of the tare load is to straighten out any initial kinks, bends, or slack in the test wire before the measurement begins. A wire that has just been taken off a spool and clamped at both ends is almost never perfectly straight, and when a small load is first applied, most of that load is used to straighten the wire rather than to stretch it. The apparent extension in the first few increments of load is therefore much larger than Hooke's law would predict for a wire that is only being stretched, because the wire is simultaneously unbending.

If the tare load were omitted and the experiment were started from the bare wire, the early data points on a load-extension graph would lie well above the true Hooke's-law line, because the extensions recorded there would include both the stretching and the unbending. The early part of the graph would curve away from the straight line expected from Hooke's law, and only after several increments, when the wire was finally straight, would the graph settle into the correct linear behaviour. The gradient extracted from such a graph would be contaminated by the curved early section and would not correspond to the true spring constant of the wire, so that Young's modulus calculated from it would be too small by a significant amount.

The tare load solves this problem by doing the unbending in advance, before any readings are taken. Once the wire is straight, every additional load produces a genuine elastic extension, and the load-extension graph is a clean straight line from the very first data point. The extension caused by the tare load itself is not part of the measurement and does not need to be known, because the zero of the micrometer is defined with the tare load already in place. Everything that the experiment records is the extension beyond that zero, and everything that the calculation uses is the slope of the straight line, both of which are independent of the unknown initial extension caused by the tare load itself.

Making Sense of the Answer: *Kipanga's suspicion was not foolish, and it is worth honouring rather than dismissing. A careful student who looks at the apparatus and asks why a weight is hanging on the wire before any measurements are taken is behaving exactly as a physicist should. The resolution is that the tare load is not part of the measurement but part of the preparation, like cleaning a test tube before pouring a chemical into it. The measurement begins only after the preparation is complete, and the zero of the measurement is the state of the wire at the end of the preparation, not the state of the wire before it.*

Think Like a Physicist: *Every measurement in physics requires a zero, and the choice of zero is almost always more delicate than it looks. In Searle's experiment the zero is the position of the micrometer after the tare load has straightened the wire. In a thermometer the zero is the reading at a well-defined reference temperature. In an electrical measurement the zero is the reading with the circuit open. Whenever you encounter a new experimental apparatus, ask yourself first: what defines the zero, and what preparation has to be done before the zero can honestly be called zero? The answer to that question often contains more physics than the rest of the procedure put together.*

HOT Example 24

A student performs Searle's experiment with the following measurements and their uncertainties: the length of the test wire is (2.50 ± 0.01)m, the diameter of the wire is (0.92 ± 0.03)mm measured with a vernier caliper, and the extension produced by a load of 50N (known exactly) is (0.90 ± 0.05)mm. Calculate (a) Young's modulus of the wire without uncertainty (b) the percentage uncertainty in Young's modulus and its absolute uncertainty (giving the final result in the standard notation), and (c) identify which of the three measured quantities contributes the most to the final uncertainty, and explain why.

Solution

(a) Young's modulus is given by Young's-modulus formula rearranged so that the area of the wire is written in terms of its diameter:

$$E = \frac{FL}{A \times \Delta L} = \frac{4FL}{\pi d^2 \times \Delta L}$$

Substituting the measured values in SI units: with the diameter squared as $(0.92 \times 10^{-3}\ \text{m})^2 = 8.464 \times 10^{-7}\ \text{m}^2$, and computing step by step:

$$E = \frac{4 \times 50\text{N} \times 2.50\text{m}}{\pi \times (0.92 \times 10^{-3}\,\text{m})^2 \times 0.90 \times 10^{-3}\text{m}} = 2.09 \times 10^{11}\ \text{Pa}$$

Therefore, the value of Young's modulus, before the uncertainty is attached, is 2.09×10^{11} Pa.

(b) Since the diameter enters the formula squared, its fractional uncertainty is doubled before it is added to the sum. The force is known exactly, so its contribution is zero:

$$\frac{\Delta E}{E} = \frac{\Delta L}{L} + \frac{\Delta(\Delta L)}{\Delta L} + 2 \times \frac{\Delta d}{d}$$

Substituting the individual fractional uncertainties:

$$\frac{\Delta E}{E} = \frac{0.01}{2.50} + \frac{0.05}{0.90} + 2 \times \frac{0.03}{0.92} = 0.125 = 12.5\%$$

The absolute uncertainty in Young's modulus is therefore:

$$\Delta E = 0.125 \times 2.09 \times 10^{11}\ \text{Pa} \approx 0.26 \times 10^{11}\ \text{Pa}$$

Rounding the value and its uncertainty to the same decimal place gives the final result:

$$E = (2.1 \pm 0.3) \times 10^{11}\ \text{Pa}$$

(c) The three fractional contributions to the total uncertainty are 0.4 percent from the length, 5.6 percent from the extension, and 6.5 percent from the diameter (after doubling). The diameter is therefore the largest contributor, even though its raw fractional uncertainty of 3.3 percent was smaller than the 5.6 percent uncertainty in the extension. The reason is the factor of two that comes from squaring the diameter in the formula for the area of the wire.

Making Sense of the Answer: *The final answer of $(2.1 \pm 0.3) \times 10^{11}$ Pa is consistent with the standard value of Young's modulus for steel, which is about 2.0×10^{11} Pa.*

Think Like a Physicist: *When a quantity in a formula is raised to a power, its fractional uncertainty is multiplied by that power before being added to the total. A quantity raised to the second power therefore contributes twice as heavily as a quantity that appears only once, and a quantity raised to the third power contributes three times as heavily. Whenever you design an experiment, identify the quantity with the highest power in the formula and spend your measurement effort on it first. A small improvement there is worth a much larger improvement in any of the linear quantities, and often a small improvement in one powered quantity is enough to turn an embarrassing experiment into a respectable one.*

Searle's apparatus has shown us how the most quoted material constant in engineering is actually extracted from a real wire in a real laboratory. But wires in real laboratories and real bridges do not only experience loads. They also experience temperature changes, and a wire that is heated expands by an amount that is usually small but is never zero. If the wire is free to expand, the expansion is harmless. If the wire is clamped between rigid supports and cannot expand, the temperature change produces a stress that can be large enough to buckle a railway line or crack a bridge beam. The next part of this chapter shows how to calculate the expansion produced by a given temperature change, the stress produced when the expansion is prevented, and what both calculations mean in the molecular picture of the potential energy curve with which this chapter began.

THERMAL EXPANSION

Linear expansion and the coefficient α

Almost every solid expands when it is heated and contracts when it is cooled. The effect is small enough to be invisible at the scale of a wristwatch but large enough at the scale of a railway bridge to rip the bridge apart if the engineer has not designed for it. For a rod of initial length L_0 that is heated through a temperature change ΔT, the increase in length ΔL is found experimentally to be proportional both to the original length and to the temperature change, and the constant of proportionality is a property of the material called its coefficient of linear expansion, α:

$$\Delta L = L_0\, \alpha\, \Delta T$$

The coefficient α has the units of reciprocal temperature and is quoted in either K^{-1} or $°C^{-1}$, which are numerically identical because a change of one kelvin equals a change of one degree Celsius.

The molecular origin of thermal expansion is already familiar from the very first section of this chapter. The potential energy curve between two atoms is not a symmetric parabola around its minimum but is steeper on the repulsive side than on the attractive side. As the temperature rises, the atoms vibrate with greater amplitude, and because the curve is asymmetric, the mean separation between atoms grows a little with every rise in temperature. The whole solid therefore becomes a little larger.

Area and volume expansion

When a flat plate is heated, it does not only become longer; every dimension in its plane grows, and the area grows with it. When a solid cube is heated, every edge grows and the volume grows. We do not need any new measurements to find the area and volume coefficients of expansion. We only need to apply the linear formula to the sides of the shape and keep the dominant terms.

Consider a square plate of side L_0 at the initial temperature. After heating through ΔT, each side has become longer by the amount $L_0\,\alpha\Delta T$, so the new side length is:

$$L = L_0 + L_0\,\alpha\Delta T = L_0(1 + \alpha\Delta T)$$

The new area of the plate is therefore the square of this new side:

$$A = L^2 = (L_0(1 + \alpha\Delta T))^2 = L_0^2\,(1 + \alpha\Delta T)^2$$

Expanding the square $(1 + \alpha\Delta T)^2$ using the identity $(a + b)^2 = a^2 + 2ab + b^2$:

$$(1 + \alpha\Delta T)^2 = 1 + 2(\alpha\Delta T) + (\alpha\Delta T)^2$$

Now comes the key step. The product $\alpha\Delta T$ is an extremely small number for the temperature changes we meet in the laboratory or in ordinary engineering. For a steel bar heated by 50K, for instance, $\alpha\Delta T = (12 \times 10^{-6}\text{K}^{-1})(50\text{K}) = 6 \times 10^{-4}$, which is a tiny fraction. Squaring such a small number gives something even tinier: $(\alpha\Delta T)^2 \approx 3.6 \times 10^{-7}$, which is about three thousand times smaller than $2(\alpha\Delta T)$ itself. To the precision of any realistic measurement we can therefore drop the $(\alpha\Delta T)^2$ term, keeping only the first-order term:

$$(1 + \alpha\Delta T)^2 \approx 1 + 2(\alpha\Delta T)$$

Substituting this back into the expression for the new area:

$$A = L_0^2(1 + 2(\alpha\Delta T)) = L_0^2 + 2L_0^2(\alpha\Delta T)$$

The first term on the right, L_0^2, is the original area A_0. The second term is therefore the change in area:

$$\boldsymbol{\Delta A = A - A_0 = 2A_0(\alpha\Delta T)}$$

This has the same form as the linear expansion formula, with a new coefficient that is exactly twice the linear coefficient. We call this new coefficient the **coefficient of area** (or **superficial**) expansion and give it the symbol $\boldsymbol{\beta}$:

$$\boldsymbol{\Delta A = A_0\,\beta\Delta T}, \text{ where } \boldsymbol{\beta = 2\alpha}$$

The same reasoning applied to a cube of side L_0 gives the volume expansion. Each side becomes $L_0(1 + \alpha\Delta T)$ after heating, so the new volume is:

$$V = L^3 = (L_0(1 + \alpha\Delta T))^3 = L_0^3\,(1 + \alpha\Delta T)^3$$

Expanding the cube using the binomial identity $(1 + x)^3 = 1 + 3x + 3x^2 + x^3$ with $x = \alpha\Delta T$:

$$(1 + \alpha\Delta T)^3 = 1 + 3(\alpha\Delta T) + 3(\alpha\Delta T)^2 + (\alpha\Delta T)^3$$

The terms $3(\alpha\Delta T)^2$ and $(\alpha\Delta T)^3$ are negligibly small next to $3(\alpha\Delta T)$ for the same reason as before, so we drop them and keep only the first-order term:

$$(1 + \alpha\Delta T)^3 \approx 1 + 3(\alpha\Delta T)$$

Substituting and simplifying:

$$V = L_0^3(1 + 3(\alpha\Delta T)) = V_0 + 3V_0(\alpha\Delta T)$$

$$\boldsymbol{\Delta V = 3V_0(\alpha\Delta T) = V_0\,\gamma\Delta T}, \text{ where } \boldsymbol{\gamma = 3\alpha}$$

The two relations $\beta = 2\alpha$ and $\gamma = 3\alpha$ are worth memorising, not as isolated facts but as direct consequences of the linear formula. If you ever forget them in an examination, you can recover them in half a minute by applying the linear formula to a square or a cube and dropping the small terms, just as we have done here.

Thermal stress when expansion is prevented

If a rod is free to expand, its length changes according to the linear formula and no stress develops inside it. But suppose the rod is clamped rigidly between two walls and then heated. The rod wants to grow longer, but the walls will not allow it to. *What happens inside the rod?*

The argument has two-steps, and it is worth following slowly.

Step one: imagine for a moment that the walls are not there. The rod would expand freely by the amount:

$$\Delta L_{free} = L_0 \, \alpha \Delta T$$

Step two: now put the walls back. The walls must somehow force the rod to stay at its original length L_0, even though the rod wanted to grow by ΔL_{free}. The only way the walls can do this is by applying an inward compressive force that squeezes the rod back by exactly the amount it wanted to grow. The compression produced by this force is given by the usual Young's modulus formula, rearranged for the change in length:

$$\Delta L_{comp} = \frac{F L_0}{AE}$$

For the rod to fit between the walls, the compression must exactly cancel the would-be expansion, so $\Delta L_{comp} = \Delta L_{free}$:

$$\frac{F L_0}{AE} = L_0 \, \alpha \Delta T$$

The original length L_0 appears on both sides of this equation and cancels. Dividing both sides by L_0 gives:

$$\frac{F}{AE} = \alpha \Delta T$$

Multiplying both sides by E and recognising that F/A is the thermal stress σ in the rod:

$$\boldsymbol{\sigma = \frac{F}{A} = E\alpha\Delta T}$$

This is the thermal stress formula. It is compressive (negative in the sign convention that calls tension positive) and has the striking property that the length of the rod has completely disappeared from the answer. A steel rod clamped between rigid walls develops exactly the same thermal stress for a given temperature change whether the rod is one centimetre long or one hundred metres long, because both the wanted expansion and the strain that results from preventing it are proportional to L_0, so the length simply cancels.

Composite rods and differential expansion

Two rods of different materials placed side by side and heated through the same temperature change expand by different amounts, because their coefficients of linear expansion differ. Two particular problems involving this differential expansion appear frequently in examinations and deserve careful derivation.

1) The constant-length-difference problem

Two rods, one of steel and one of copper, are placed side by side. At a reference temperature, the steel rod has length L_s and the copper rod has length L_c. We want the difference between their lengths to remain exactly the same at every temperature, so that the end of the steel rod is always a fixed distance away from the end of the copper rod. When the temperature changes by ΔT, the steel rod expands by $L_s \, \alpha_s \, \Delta T$ and the copper rod expands by $L_c \, \alpha_c \, \Delta T$. For the difference to remain constant, these two expansions must be exactly equal:

$$L_s \, \alpha_s \, \Delta T = L_c \, \alpha_c \, \Delta T$$

The common factor ΔT cancels from both sides, leaving a condition that involves only the lengths and the coefficients:

$$L_s \, \alpha_s = L_c \, \alpha_c$$

Rearranging for the ratio of the two lengths:

$$\frac{L_s}{L_c} = \frac{\alpha_c}{\alpha_s}$$

This is the key result. The two rod lengths must be in the inverse ratio of their coefficients of expansion: the rod with the smaller coefficient (steel) must be made longer, and the rod with the larger coefficient (copper) must be made shorter, so that they expand by equal amounts. This is exactly the working principle of the compensated clock pendulum, whose precision depends on the gap between its two supporting rods staying constant through all seasons of the year.

2) The clamped-composite problem

Two rods of different materials are joined end to end and the whole assembly is clamped rigidly between two walls. When the assembly is heated by ΔT, neither rod can expand, because the walls hold the two ends fixed. Our task is to find the thermal stress developed in the composite. Assume that the two rods have the same cross-sectional area A. We build the argument in four short steps.

Step 1. Compute the free expansion. If the walls were absent, each rod would expand freely according to the linear formula. The total free expansion of the composite is the sum of the two individual expansions:

$$\Delta L_{free} = L_1 \, \alpha_1 \, \Delta T + L_2 \, \alpha_2 \, \Delta T$$

Step 2. Recognise that the walls transmit the same force through both rods. A compressive force F applied by the walls at the outer ends must also act at the joint between the two rods, because the joint is in equilibrium. Since both rods have the same cross-sectional area A, the stress $\sigma = F/A$ is the same in both rods, even though the strain in each is different.

Step 3. Compute the compression of each rod under this stress. Using Young's modulus for each material, the strain in rod 1 is σ/E_1 and the strain in rod 2 is σ/E_2. The compression of each rod is its strain times its length:

$$\Delta L_1 = \left(\frac{\sigma}{E_1}\right) L_1 = \frac{\sigma L_1}{E_1}$$
$$\Delta L_2 = \left(\frac{\sigma}{E_2}\right) L_2 = \frac{\sigma L_2}{E_2}$$

Step 4. Set the total compression equal to the total free expansion, because the walls force the net change in length to be zero:

$$\frac{\sigma L_1}{E_1} + \frac{\sigma L_2}{E_2} = L_1\,\alpha_1\,\Delta T + L_2\,\alpha_2\,\Delta T$$

Factoring out σ on the left and solving for it:

$$\boldsymbol{\sigma = \frac{(L_1\,\alpha_1 + L_2\,\alpha_2)\Delta T}{\frac{L_1}{E_1} + \frac{L_2}{E_2}}}$$

As a consistency check, set both materials equal: $E_1 = E_2 = E$ and $\alpha_1 = \alpha_2 = \alpha$. The numerator becomes $(L_1 + L_2)\alpha\Delta T$ and the denominator becomes $(L_1 + L_2)/E$, so the stress reduces to $\sigma = E\alpha\Delta T$, the single-material formula from the previous subsection. The composite formula is therefore a clean generalisation of the simpler result.

Bimetallic strips and applications

A bimetallic strip is a thin bar formed by bonding two layers of different metals along their entire length. When the strip is heated, the metal with the larger coefficient of expansion tries to grow more than the metal with the smaller coefficient, but the bond between them prevents either layer from sliding over the other. The strip therefore bends into an arc, with the metal of larger α on the outside of the curve (because the outside of a curve is longer than the inside) and the metal of smaller α on the inside. On cooling below the reference temperature, the strip bends the other way. This simple device is the heart of almost every mechanical thermostat: the bending strip opens or closes an electrical contact as the temperature crosses a set point, and by choosing the two metals and the thickness of the strip, the designer can tune the set point over a wide range.

Thermal expansion shows up in many other pieces of everyday engineering. Expansion joints are built into long railway lines, concrete slab roads, and bridges to give the material somewhere to go when it grows. The Selander Bridge in Dar es Salaam uses finger joints between its concrete deck spans for exactly this reason. Hot rivets were historically driven through holes in steel plates and allowed to cool in place, so that as they contracted they pulled the plates tightly together with an enormous clamping force. Shrink-fitting exploits the same effect: a steel ring is heated until it expands enough to slip over a shaft, then cools and grips the shaft so tightly that it can transmit large torques without keys or welds. Railway wheels, turbine discs, and gear rings are all assembled by shrink-fitting.

Ideas, like rails, expand when heated; and these ones are ready to buckle. Let us clamp them down with a few worked examples.

BINDER Example 25

A length of steel railway rail is installed with its two ends rigidly clamped so that it cannot change its length. The rail is laid at $20°C$ and the ambient temperature then rises to $65°C$ during the hottest part of the afternoon. Taking the Young's modulus of steel as $2 \times 10^{11}\text{Pa}$ and its coefficient of linear expansion as $12 \times 10^{-6}\text{K}^{-1}$, calculate the compressive thermal stress developed in the rail.

Solution

The temperature change is $\Delta T = 65°C - 20°C = 45\text{K}$. Because the rail cannot change its length, the thermal expansion that would naturally occur is cancelled by an equal compressive strain, and the resulting thermal stress is given by the formula $\sigma = E\alpha\Delta T$. Substituting:

$$\sigma = E\alpha\Delta T = (2 \times 10^{11}\text{Pa})(12 \times 10^{-6}\text{K}^{-1})(45\text{K}) = 1.08 \times 10^{8}\text{Pa}$$

The thermal stress developed in the clamped rail is therefore $1.08 \times 10^{8}\text{Pa}$, (or 108MPa).

Making Sense of the Answer: *The answer depends only on the Young's modulus, the coefficient of linear expansion, and the temperature change. Notice that the length of the rail has vanished from the calculation: a clamped steel rail develops the same thermal stress of 108MPa for a 45K temperature rise whether the rail is a centimetre long or a kilometre long. What changes with length is not the stress but the tendency to buckle sideways when that stress is applied.*

Think Like a Physicist: *Whenever a thermal stress calculation gives a large number, ask whether the actual failure mode will match what the calculation assumed. Direct crushing is only one way a clamped rod can give way. Lateral buckling, joint slippage, and fracture at pre-existing flaws can all occur at stresses far below the raw compressive strength, especially in long slender pieces. A stress well below the crushing limit is not a guarantee of safety, only a single box that has been ticked.*

REAL Example 26

Expansion gaps are built into every long railway line and into the concrete slab roads of the Morogoro highway. Explain in physical terms why these gaps are necessary, what would go wrong if they were omitted, and why, despite their presence, the

gaps do not produce a noticeably bumpy ride for trains or for vehicles travelling along the road at normal ambient temperatures.

Solution

A continuous rail or continuous concrete slab that is rigidly held at its ends cannot expand freely when the temperature rises, and instead develops a compressive thermal stress of magnitude $\sigma = E\alpha\Delta T$. For steel rails heated from a cool morning temperature of $20°C$ to an afternoon temperature of $60°C$, this thermal stress is of the order of 10^8Pa, which is not enough to crush the steel but is easily enough to cause lateral buckling of a slender rail fixed to its sleepers. A buckled rail twists sideways into a sinuous shape and can derail a train. Concrete slabs subjected to the same kind of thermal stress crack in compression during the hottest part of the day, and can also crack in tension on cold nights if they are bonded to neighbouring slabs, because concrete is much weaker in tension than in compression.

Expansion gaps cure both problems by giving the material somewhere to grow. When the rail warms up, it expands into the gap and no thermal stress develops. The size of the gap is chosen so that it just closes at the highest temperature ever expected on that stretch of line, with a small safety margin. A gap that is too small will still allow the rail to buckle on the hottest day; a gap that is too large will widen uncomfortably on cold nights, when the rail contracts.

The gaps do not cause a bumpy ride under normal conditions because they are very small compared with the size of a train wheel or a vehicle tyre. A typical expansion gap on a railway line is only a few millimetres wide at normal temperatures, and the contact patch between the wheel and the rail is larger than this, so the wheel always has most of itself resting on solid rail as it rolls across the gap. A small but audible click is sometimes heard when the wheels cross each joint, especially at low temperatures when the gap is wider, but the jolt is too small to feel in a moving carriage. Concrete slab roads use the same principle: the gaps are thin enough that tyres bridge them easily, and the flexible sealing compound that fills them further smooths the transition.

Making Sense of the Answer: *The design of an expansion gap is a quiet compromise between two different mechanical failures. Without any gap at all, the material fails in compression at high temperature. With too large a gap, the wheels begin to jolt and the joints themselves become damaged by repeated impact. The correct size is the one that just closes on the hottest expected day and never opens dangerously wide on the coldest expected night.*

Think Like a Physicist: *Whenever a civil engineering structure carries a large load over a long distance, look for the thermal compromise hidden somewhere in its geometry. Bridges have expansion joints at their ends or in the middle. Railway lines have gaps between rails. Long pipelines have zigzag bends or coils. Power lines sag deliberately so that they can contract on cold nights without snapping. In every case, the same physical question is being answered: where will the thermal expansion go, and how can we let it go there without damaging anything?*

HOT Example 27

At a temperature of $20°C$, a steel rod is exactly 5.0cm longer than a copper rod placed beside it. The rods are to be used in a precision instrument in which the difference in their lengths must remain constant at every temperature. Given that $\alpha_s = 12 \times 10^{-6}\text{K}^{-1}$ and $\alpha_c = 17 \times 10^{-6}\text{K}^{-1}$, calculate the actual length of each rod at $20°C$.

Solution

For the difference in length to remain the same at every temperature, both rods must expand by exactly the same absolute amount whenever they are heated by the same ΔT. Setting the two linear expansions equal:

$$L_s\,\alpha_s\,\Delta T = L_c\,\alpha_c\,\Delta T$$

The common factor ΔT cancels from both sides, leaving:

$$L_s\,\alpha_s = L_c\,\alpha_c$$

Substituting the numerical values of the coefficients:

$$L_s \times 12 = L_c \times 17$$

Rearranging gives the length ratio $\frac{L_s}{L_c} = \frac{17}{12}$. The second condition given in the question, that the steel rod is 5.0cm longer than the copper rod, provides a second equation:

$$L_s - L_c = 5.0\text{cm}$$

Substituting $L_s = (\frac{17}{12})\,L_c$ from the first equation into the second:

$$\left(\frac{17}{12}\right)L_c - L_c = 5.0\text{cm}$$

Solving for L_c:

$$L_c = 12\text{cm}$$

And therefore $L_s = L_c + 5.0\text{cm} = 12\text{cm} + 5.0\text{cm} = 17\text{cm}$.
The copper rod is 12cm long and the steel rod is 17cm long at $20°C$.

Making Sense of the Answer: *Both rods expand when the instrument is heated, but because the copper rod is shorter and the steel rod is longer in just the right proportion, they expand by identical absolute amounts. The 5.0cm gap between their*

free ends therefore remains exactly 5.0cm at every temperature. The particular length difference of 5.0cm was incidental; any desired gap can be arranged by scaling both rods up or down while keeping their length ratio at 17:12.

Think Like a Physicist: *The general lesson of this problem is that whenever two effects must cancel at every temperature, it is their rates of change that must match, not their absolute values. Setting rates equal at the outset converts a physics problem about expansion into an algebra problem about ratios, and the algebra is usually easier than the physics.*

HOT Example 28

A solid steel cube of side 5.0cm at 20°C is held inside a perfectly rigid steel jig that prevents the cube from expanding in any direction. The cube is then heated uniformly to 70°C. The bulk modulus of steel is 1.6×10^{11}Pa and its coefficient of linear expansion is $12 \times 10^{-6}\text{K}^{-1}$. (a) Derive an expression for the hydrostatic compressive pressure that the jig must exert on the cube to prevent it from expanding, in terms of the bulk modulus B, the coefficient of linear expansion α, and the temperature change ΔT. (b) Calculate the numerical value of this pressure. (c) Calculate the total inward force exerted on each face of the cube by the jig.

Solution

(a) The temperature change is $\Delta\text{T} = 70°\text{C} - 20°\text{C} = 50\text{K}$. The argument has two-steps, exactly parallel to the linear case but applied in three dimensions instead of one.

Step 1: Find the volume the cube would occupy if it were free to expand

If the jig were not there, the cube would expand in all three directions according to the volume coefficient of expansion $\gamma = 3\alpha$. The fractional change in volume would be:

$$\frac{\Delta\text{V}_{\text{free}}}{\text{V}_0} = \gamma\Delta\text{T} = 3\alpha\Delta\text{T}$$

Step 2: Find the pressure required to compress the cube back to its original volume

The jig holds the cube at exactly its original volume V_0, so the cube cannot grow at all. The fractional volume change must therefore be cancelled by an equal and opposite fractional volume change produced by an inward hydrostatic pressure. From the definition of the bulk modulus:

$$\text{B} = -\frac{\Delta\text{P}}{\frac{\Delta\text{V}}{\text{V}}}$$

The pressure required to compress the cube by a fractional volume of $-\Delta\text{V}_{\text{free}}/\text{V}_0 = -3\alpha\Delta\text{T}$ is therefore:

$$\Delta\text{P} = -\text{B} \times \frac{\Delta\text{V}}{\text{V}_0} = -\text{B}(-3\alpha\Delta\text{T}) = 3\text{B}\alpha\Delta\text{T}$$

So the hydrostatic compressive pressure exerted by the jig on the cube is:

$$\boxed{\mathbf{\Delta P = 3B\alpha\Delta T = B\gamma\Delta T}}$$

This is the three-dimensional analogue of the linear thermal stress formula $\sigma = \text{E}\alpha\Delta\text{T}$ that appeared in the example 25. The Young's modulus E has been replaced by the bulk modulus B, and the linear coefficient α has been replaced by the volume coefficient $\gamma = 3\alpha$, which is exactly what we should expect when the geometry generalises from one axis to three.

(b) Substituting the numerical values with units carried at every step:

$$\Delta\text{P} = 3(1.6 \times 10^{11}\text{Pa})\left(12 \times 10^{-6}\text{K}^{-1}\right)(50\text{K}) = 2.88 \times 10^{8}\text{Pa}$$

The hydrostatic pressure that the jig must exert on the cube is therefore approximately 2.88×10^{8}Pa, or about 288MPa.

(c) The area of one face of the cube is $\text{A} = (5.0 \times 10^{-2}\text{m})^2 = 2.5 \times 10^{-3}\text{m}^2$.

The total inward force exerted on each face is the pressure multiplied by the area:

$$\text{F} = \Delta\text{P} \times \text{A} = (2.88 \times 10^{8}\text{Pa})(2.5 \times 10^{-3}\text{m}^2) = 7.2 \times 10^{5}\text{N}$$

So each of the six faces of the cube is being pushed inward by a force of approximately 7.2×10^{5}N.

Making Sense of the Answer: *The pressure of* 288MPa *is about three times larger than the* 108MPa *thermal stress we found for a clamped steel rail in the linear case under the same temperature rise. The factor of three is not a coincidence: it appears because the cube is being prevented from expanding in three independent directions, not just one. The bulk modulus* B *relates a hydrostatic pressure to a fractional volume change, and the volume change for the same temperature rise is three times larger than the length change, so the pressure required to cancel it is correspondingly larger. The formula* $\Delta\text{P} = 3\text{B}\alpha\Delta\text{T}$ *encodes both effects at once.*

Think Like a Physicist: *When a solid is constrained against expansion, the right modulus to use depends on which dimensions are being held fixed. If only one length is held fixed (a clamped rod), use Young's modulus and the linear coefficient:* $\sigma = \text{E}\alpha\Delta\text{T}$. *If all three dimensions are held fixed (a clamped block, or a fluid sealed inside a rigid container), use the bulk modulus and the volume coefficient:* $\Delta\text{P} = \text{B}\gamma\Delta\text{T} = 3\text{B}\alpha\Delta\text{T}$. *The two formulas have the same algebraic shape because the underlying physics is the same because prevented expansion produces stress proportional to the modulus that resists that*

kind of deformation. Whenever a thermal stress problem appears, the first question to ask is how many dimensions of the solid are constrained. The answer tells you which modulus to reach for.

HOT Example 29

A composite rod is made by joining a copper rod of length 30cm to a steel rod of length 20cm end to end. The two rods have the same cross-sectional area, and the composite rod is clamped rigidly between two walls at 20°C. The whole assembly is then heated to 70°C. Calculate the thermal stress developed in the composite, given that for copper $E_c = 1.2 \times 10^{11}\text{Pa}$ and $\alpha_c = 17 \times 10^{-6}\text{K}^{-1}$, and for steel $E_s = 2.0 \times 10^{11}\text{Pa}$ and $\alpha_s = 12 \times 10^{-6}\text{K}^{-1}$.

Solution

The temperature change is $\Delta T = 70°\text{C} - 20°\text{C} = 50\text{K}$.

The total free expansion that would occur if the walls were absent is the sum of the free expansions of the copper and steel sections:

$$\Delta L_{\text{free}} = L_c\,\alpha_c\,\Delta T + L_s\,\alpha_s\,\Delta T$$

Substituting the numerical values with units:

$$\Delta L_{\text{free}} = (0.30\text{m})(17 \times 10^{-6}\text{K}^{-1})(50\text{K}) + (0.20\text{m})(12 \times 10^{-6}\text{K}^{-1})(50\text{K}) = 3.75 \times 10^{-4}\text{m}$$

The walls transmit the same compressive force through both rods, so the stress $\sigma = F/A$ is the same in both, even though the strain in each is different. So the compression of each rod, using Young's modulus:

$$\Delta L_c = \frac{\sigma L_c}{E_c}, \quad \Delta L_s = \frac{\sigma L_s}{E_s}$$

The total compression must equal the total free expansion, because the walls hold the ends of the composite fixed:

$$\frac{\sigma L_c}{E_c} + \frac{\sigma L_s}{E_s} = \Delta L_{\text{free}}$$

Making σ the subject:

$$\sigma = \frac{\Delta L_{\text{free}}}{\frac{L_c}{E_c} + \frac{L_s}{E_s}}$$

Substituting:

$$\sigma = \frac{3.75 \times 10^{-4}\text{m}}{\frac{0.30\text{m}}{1.2 \times 10^{11}\text{Pa}} + \frac{0.20\text{m}}{2.0 \times 10^{11}\text{Pa}}} = 1.07 \times 10^{8}\text{Pa}$$

The thermal stress in the composite rod is therefore approximately $1.07 \times 10^{8}\text{Pa}$, (or about 107MPa), compressive throughout both sections.

Making Sense of the Answer: *The thermal stress in this two-material composite is almost identical to the 108MPa we found for a single clamped steel rail in Example 25. This is not a coincidence. Copper expands more than steel but has a smaller Young's modulus, so its extra expansion is roughly cancelled by its greater willingness to be compressed. The simpler single-material formula $\sigma = E\alpha\Delta T$ is therefore a reasonable first approximation for any composite whose components are not wildly different, and the more careful composite formula refines that first approximation rather than overturning it.*

Think Like a Physicist: *The key step in composite-rod problems is recognising that when two pieces share the same cross-section and the same clamping force, the stress is the same in both pieces even though the strain is not. Stress is force per unit area, which is transmitted unchanged from one piece to the next across the joint. Strain is how much each material yields to that stress, which depends on its stiffness. Keeping these two quantities separate is the difference between a clean solution and a paragraph of confusion.*

Thermal expansion now joins Young's modulus, the stress-strain curve, and Searle's apparatus on the list of ideas that form the working physics of elastic solids. The last remaining task for this chapter is to step back from the derivations and calculations, gather together the engineering and biological situations in which elasticity is quietly doing its work, and ask what the whole chapter has been preparing us to understand.

APPLICATIONS OF ELASTICITY

The ideas built up in the previous sections are not abstract. Every modulus, every stress, every strain, and every coefficient of expansion has a job to do in the world outside the physics classroom. This final section of the chapter gathers those jobs together and shows how the single body of elasticity theory underwrites the design of buildings and lifts, the shape of the beams that hold up a floor, the mechanical operation of the human body, and the slow failure of springs and cables through years of use. None of these applications requires new physics. They require only the careful application of what has already been learned.

1. Safety factors in engineering design

Every engineering structure is designed to carry a maximum expected load, and every material has a stress beyond which it will yield or fracture. In a perfect world, a structure designed to the exact yield stress of its material would work correctly at

its design load and fail by only the tiniest margin. In the real world this is never done, because the real world contains hidden sources of error: material flaws beneath the surface, loads that fluctuate unpredictably, temperature changes, corrosion, manufacturing variations, and the occasional careless operator. To protect against all these sources of uncertainty, engineers divide the material's yield stress (or ultimate tensile strength) by a number called the safety factor, and use the result as the working stress that the structure must never exceed in normal service:

$$\sigma_{\text{working}} = \frac{\sigma_{\text{yield}}}{n}$$

Here n is the safety factor, typically between 2 and 10 depending on the application. For ordinary structural steel beams in a building, safety factors of 2 or 3 are common. For lift and elevator cables, which must carry people, the factor is usually between 4 and 6. For aircraft components, whose failure can be fatal and whose load history is uncertain, safety factors of 8 or more are standard, and for crane cables lifting overhead loads above workers, factors of up to 10 are required by law. The exact value chosen depends on the consequences of failure, the reliability of the load estimate, the quality of the material, and the cost of adding extra strength. A higher safety factor is always safer but always more expensive.

The procedure for using a safety factor in design is always the same. Start from the worst-case load F that the component must be able to carry. Compute the corresponding working stress σ_{working} from the yield stress and the chosen safety factor. Rearrange the definition of stress to find the minimum cross-sectional area $A = \frac{F}{\sigma_{\text{working}}}$ needed to keep the stress at or below the working value. If the component is a circular wire or cable, convert the area to a diameter using $d = 2\sqrt{\frac{A}{\pi}}$. The result is the smallest diameter that will survive the worst-case load with the chosen margin of safety.

2. Structural geometry and the I-beam

The stress inside a loaded beam is not uniform across its cross-section. When a horizontal beam bends under a downward load, the material at the top of the beam is compressed, the material at the bottom is stretched, and the material along the middle horizontal line (called the neutral axis) experiences no stress at all. The stress at any point in the cross-section grows in proportion to its distance from the neutral axis, which means that the material near the top and bottom of the beam is doing most of the work of resisting the load, while the material along the middle is almost idle.

Engineers exploit this fact by removing the idle material from the middle and concentrating it at the top and bottom, where it earns its place. The result is the I-beam, whose cross-section resembles the capital letter I: two broad horizontal flanges at the top and bottom, connected by a thin vertical web in the middle. An I-beam uses much less material than a solid rectangular beam of equivalent bending strength, which is why railway rails, bridge girders, floor joists, and the frames of tall buildings are almost always I-shaped. The rectangular beam is simpler but wasteful. The I-beam is a direct application of stress analysis to the economy of structural design.

3. Biological elasticity: four tissues, four stiffnesses

The body of every animal is a collection of elastic structures, each one tuned by evolution to a specific mechanical task. Four tissues in the human body illustrate the range of Young's moduli found in living material, spanning nearly five orders of magnitude from the softest to the stiffest.

Tendons connect muscles to bones and must transmit the full force of a contracting muscle without stretching significantly, because any stretch would waste energy that the muscle meant to deliver to the bone. Tendons therefore have a high Young's modulus of about 1.5GPa, which is stiff enough to behave almost like a rigid rope under normal loads. A tendon stretching by even 4% is close to rupture, which is why tendon injuries are so common in athletes.

Arteries must do the opposite of what tendons do. They must stretch and contract with every heartbeat to smooth the pulsatile flow of blood from the ventricles into a steady flow in the smaller vessels. Arteries therefore have a Young's modulus of only about 0.5MPa, roughly three thousand times smaller than that of tendons, and strains of 10% to 15% during each cardiac cycle are normal. With age, the walls of arteries become progressively stiffer, a condition called arteriosclerosis, and the artery loses its ability to buffer the pressure pulse. This is one reason why elderly patients are more prone to high blood pressure and to the rupture of small blood vessels in the brain.

Cartilage covers the ends of bones at every joint and acts as a shock absorber during walking, running, and jumping. Its Young's modulus of about 1MPa is intermediate between tendon and artery, giving it enough firmness to bear load and enough compliance to cushion impacts. Cartilage has no blood supply of its own and therefore heals slowly once damaged, which is why knee and hip cartilage injuries often become chronic.

Bone is the stiffest biological tissue, with a Young's modulus of about 20GPa for compact cortical bone. This is as stiff as concrete and a little stiffer than wood along the grain. The high stiffness of bone allows it to form the rigid skeleton that supports the body, and its strength per unit mass is actually better than that of most metals, which is one reason a bird's hollow bones can be so light and yet carry the whole weight of the bird in flight.

Put together, the four tissues span nearly five orders of magnitude in stiffness: artery at 0.5MPa, cartilage at 1MPa, tendon at 1.5GPa, and bone at 20GPa. Each stiffness is matched to a mechanical job. A body built entirely from tendon would have no cushioning at the joints and no pressure buffering in the arteries. A body built entirely from artery-grade tissue would collapse under its own weight. Evolution has chosen the right material for each part of the body by picking the right value of Young's modulus.

4. Elastic fatigue in everyday life

Elastic fatigue is the quiet failure mode that catches engineers and students by surprise. A material that has been loaded many times at a stress well below its breaking point will eventually fail, because tiny cracks at microscopic flaws grow a little with each loading cycle until one of them becomes large enough to cause sudden fracture. The failure is not preceded by any visible warning, which is why elastic fatigue has been responsible for many historical engineering disasters and continues to determine the inspection schedules of aircraft, bridges, and pressure vessels today.

A spring balance in daily use at a market loses its zero progressively over months and years, not because the spring has ever been overloaded in any single use but because thousands of loading and unloading cycles have accumulated microscopic plastic deformation at internal crack tips. A spring balance that reads two hundred grams when empty is not broken; it is suffering from the same condition in slow motion that kills aircraft wings in years and bridge cables in decades. The difference is only the number of cycles required for the damage to become visible.

Aircraft wings flex slightly with every takeoff, landing, and gust of turbulence. The fatigue life of a wing spar is measured in flight cycles, not flight hours, and the number of cycles at which fatigue cracks become detectable is what determines how often the aircraft must be inspected and eventually retired. The first generation of commercial jet aircraft in the 1950s included the de Havilland Comet, three of which broke up in flight because stress concentrations at the corners of their square windows caused fatigue cracks to grow and propagate catastrophically. The disaster taught engineers to round off every corner and changed aircraft design forever.

Human bone is also vulnerable to elastic fatigue. A runner who suddenly increases training mileage can develop stress fractures in the tibia or metatarsals after weeks or months of cumulative microscopic bone damage, even though no single training session exceeded what the bone could withstand. The same physics that kills a bridge cable in fifty years can break an athlete's shin in three months, and the treatment is the same in both cases: stop loading, let the cracks heal, and do not repeat the pattern that caused them.

REAL Example 30

A spring balance used daily in a busy Kariakoo market reads 200g when empty, even though it read zero when it was new several years ago. The trader has never hung a load on the balance that was anywhere near its rated maximum, so the spring has never been overloaded in any single use. Explain the physical process responsible for the progressive loss of the zero reading, identify how it differs from ordinary overloading, and suggest one practical step that the trader could take to slow the rate at which the error develops.

Solution

The progressive loss of the zero reading is caused by elastic fatigue in the steel spring of the balance. Every time a load is hung on the balance and then removed, the spring is stretched and then released. On any single cycle, the stretch is well within the elastic limit of the steel, and the spring appears to return exactly to its original length. In reality, microscopic flaws on the inside and outside surfaces of the wire grow a tiny amount with each cycle. The growth is too small to see or to detect in a single measurement, but it accumulates slowly over thousands of cycles into a permanent plastic deformation of the spring, which is why the unloaded length of the spring gradually grows longer and the pointer slowly drifts away from zero.

This is completely different from the damage caused by ordinary overloading. Overloading means stretching the spring beyond its elastic limit in a single event, which produces an immediate and obvious permanent deformation. Elastic fatigue, by contrast, produces no damage that any single measurement can detect, and is caused by repeated loading well within the elastic limit. Elastic fatigue damage accumulates cycle by cycle over a long period of time, while overloading damage happens all at once. The two processes produce similar-looking results at the end, but they arise from very different causes, and the first is much harder to guard against than the second.

To slow the rate at which the zero error develops, the trader could remove the load from the balance whenever it is not in active use, instead of leaving heavy goods hanging on the hook between customers. Each hour that the spring sits under load is an hour in which the crack tips inside the metal are being stressed, even though no measurement is being taken. Unloading the balance between uses reduces the total number of stress-hours the spring experiences over its life and extends the time before fatigue becomes significant. A second step, just as valuable, is to recalibrate the balance regularly against a known standard mass and to apply a correction to every reading until the day comes when the error becomes too large to correct, at which point the spring must be replaced.

Making Sense of the Answer: *The steady growth of the zero error is a direct analogue of the cracks that grow in an aircraft wing spar or a bridge cable. Markets, aircraft, and bridges all fail by the same mechanism, separated only by the number of*

cycles required for the damage to become visible. The spring balance has a short fatigue life because it is cycled many times each day; the bridge cable has a long fatigue life because it is cycled only a few times per minute as traffic passes over it. Both are on the same clock, just set to different speeds.

Think Like a Physicist: *Whenever an instrument is said to be 'in need of calibration' after long use, ask whether the underlying cause is elastic fatigue rather than damage. Elastic fatigue cannot be prevented by any amount of careful use; it can only be slowed down by reducing the number of loading cycles, reducing the amplitude of each cycle, or replacing the component before the damage becomes catastrophic. The trader's habit of unloading the balance between customers is in the first category. The inspection schedule of an aircraft fleet is in the third. Both are expressions of the same underlying idea: material fatigue is something you manage, not something you cure.*

HOT Example 31

A passenger lift is to be designed to carry a maximum live load of 600kg (including passengers and goods), supported by four identical parallel steel cables. The yield stress of the steel used for the cables is $2.5 \times 10^8\text{Pa}$, and the safety factor required by the local building code is 4. Taking $g = 9.8\text{m/s}^2$, calculate (a) the total weight of the maximum live load, (b) the maximum force that each cable must support, (c) the maximum working stress allowed in each cable, and (d) the minimum diameter of each cable.

Solution

(a) The total weight W of the maximum live load is the mass times the gravitational field strength:

$$W = mg = (600\text{kg})(9.8\text{m/s}^2) = 5880\text{N}$$

The maximum live load therefore exerts a downward weight of 5880N on the four cables combined.

(b) The four cables share the load equally, so the force on each one is one quarter of the total:

$$F = \frac{W}{4} = \frac{5880\text{N}}{4} = 1470\text{N}$$

Each cable must therefore support a force of 1470N in normal service with the maximum live load.

(c) The maximum working stress is the yield stress divided by the safety factor:

$$\sigma_{\text{working}} = \frac{\sigma_{\text{yield}}}{n} = \frac{2.5 \times 10^8\text{Pa}}{4} = 6.25 \times 10^7\text{Pa}$$

Each cable is therefore allowed to experience a stress of at most $6.25 \times 10^7\text{Pa}$ in normal service. This is the stress that corresponds to the full live load divided by four and still leaves a factor of four in hand before the steel begins to yield.

(d) The minimum cross-sectional area of each cable is the force on the cable divided by the maximum allowed stress:

$$A = \frac{F}{\sigma_{\text{working}}} = \frac{1470\text{N}}{6.25 \times 10^7\text{Pa}} = 2.352 \times 10^{-5}\text{m}^2$$

Each cable is circular, so its cross-sectional area is related to its diameter by $A = \pi d^2/4$. Rearranging for the diameter:

$$d = \sqrt{\frac{4A}{\pi}} = \sqrt{\frac{4 \times 2.352 \times 10^{-5}\text{m}^2}{\pi}} = 5.47 \times 10^{-3}\text{m}$$

Therefore each cable must have a diameter of at least $5.47 \times 10^{-3}\text{m}$ (or about 5.5mm), to carry the maximum live load with the required safety factor of four.

Making Sense of the Answer: *A diameter of 5.5mm sounds thin for a lift cable, and in fact real lift cables are almost always thicker than this; typically, 8mm to 12mm for a small passenger lift. The difference comes from three considerations that the simple calculation above did not include. First, the mass of the lift car itself is additional dead weight that has to be supported in addition to the live load of passengers. Second, dynamic effects when the lift starts, stops, or changes direction add extra forces on top of the static weight, often as much as 50 percent more. Third, modern lifts use counterweights to reduce the net force on the cables, which can partially offset the first two effects. The 5.5mm answer should therefore be understood as the absolute minimum dictated by the safety factor for the stated live load alone, not as a practical design specification.*

Think Like a Physicist: *Design-feasibility problems always have the same four-step structure: (i) find the worst-case load, (ii) divide it among the load-bearing members to find the force on each, (iii) compute the working stress from the yield stress and the safety factor, and (iv) rearrange the definition of stress to find the minimum cross-section or diameter needed. Keeping the four steps separate prevents the easy mistake of dividing by the safety factor twice or forgetting to divide at all. Whenever you are asked to size a cable, a rod, or a column to carry a given load, recognise the pattern and walk through the steps in order.*

The chapter has walked us from the PE curve to the cross-section of a lift cable, and along the way it has let stress and strain, the three moduli, Searle's apparatus, and rails buckling in the afternoon sun all introduce themselves one at a time. They will no longer be polite. The miscellaneous worked examples that follow mix Young's modulus with thermal stress, hysteresis with safety factors, and hanging wires with breaking points, the way real NECTA questions do.

MISCELLANEOUS WORKED EXAMPLES ON ELASTICITY

Example 32

(a) Distinguish between the proportional limit, the elastic limit, and the yield point of a metal rod under tensile loading.

(b) A steel rod of length 2.0m and diameter 2.0mm is subjected to an increasing tensile force. At a load of 500N the rod extends by 1.60mm. When the load is increased to 1200N the rod extends by 3.80mm, and beyond 1200N the rod begins to deform plastically. (i) Calculate Young's modulus of the steel from the first pair of readings. (ii) Calculate the stress in the rod at the elastic limit. (iii) Calculate the elastic potential energy stored in the rod at the elastic limit.

Solution

(a) The proportional limit is the highest stress up to which stress is directly proportional to strain, that is, up to which Hooke's law holds. The elastic limit is the highest stress at which the rod still returns exactly to its original length when the load is removed, and lies just above the proportional limit. The yield point is the stress at which the rod begins to flow plastically, giving a large increase in strain for a small increase in stress, and lies above the elastic limit.

(b) (i) Cross-sectional area of the rod:

$$A = \frac{\pi d^2}{4} = \frac{\pi(2.0 \times 10^{-3}\text{m})^2}{4} = 3.14 \times 10^{-6}\text{m}^2$$

Young's modulus from the first pair of readings:

$$E = \frac{FL}{A\Delta L} = \frac{(500\text{N})(2.0\text{m})}{(3.14 \times 10^{-6}\text{m}^2)(1.60 \times 10^{-3}\text{m})} = 1.99 \times 10^{11}\text{Pa}$$

(ii) Stress at the elastic limit (load 1200N):

$$\sigma = \frac{F}{A} = \frac{1200\text{N}}{3.14 \times 10^{-6}\text{m}^2} = 3.82 \times 10^{8}\text{Pa}$$

(iii) Elastic potential energy stored at the elastic limit:

$$U = \frac{1}{2}F\Delta L = \frac{1}{2}(1200\text{N})(3.80 \times 10^{-3}\text{m}) = 2.28\text{J}$$

Example 33

(a) A steel wire and a copper wire of equal length and equal cross-section are joined end to end and stretched by the same tensile force. Given that the Young's modulus of steel is greater than that of copper, explain which wire extends more, and state the physical reason.

(b) A composite wire is formed by joining a steel wire of length 2.0m to a copper wire of length 3.0m end to end. Both wires have the same cross-sectional area of $1.5 \times 10^{-6}\text{m}^2$. A tensile force of 450N is applied to the free ends. Given $E_s = 2.0 \times 10^{11}\text{Pa}$ and $E_c = 1.2 \times 10^{11}\text{Pa}$, calculate (i) the tensile stress in each wire, (ii) the extension of each section, and (iii) the total extension of the composite.

Solution

(a) The copper wire extends more. With equal cross-section and equal force, both wires carry the same stress. Strain is stress divided by Young's modulus, so the wire with the smaller Young's modulus (copper) has the larger strain. With equal lengths, the larger strain produces the larger extension.

(b) (i) The two wires are in series, so they carry the same force and have the same cross-section. The stress in both is therefore equal:

$$\sigma = \frac{F}{A} = \frac{450\text{N}}{1.5 \times 10^{-6}\text{m}^2} = 3.0 \times 10^{8}\text{Pa}$$

(ii) Extension of the steel section:

$$\Delta L_s = \frac{\sigma L_s}{E_s} = \frac{(3.0 \times 10^{8}\text{Pa})(2.0\text{m})}{2.0 \times 10^{11}\text{Pa}} = 3.0 \times 10^{-3}\text{m} = 3.0\text{mm}$$

Extension of the copper section:

$$\Delta L_c = \frac{\sigma L_c}{E_c} = \frac{(3.0 \times 10^{8}\text{Pa})(3.0\text{m})}{1.2 \times 10^{11}\text{Pa}} = 7.5 \times 10^{-3}\text{m} = 7.5\text{mm}$$

(iii) Total extension of the composite:

$$\Delta L = \Delta L_s + \Delta L_c = 3.0\text{mm} + 7.5\text{mm} = 10.5\text{mm}$$

Example 34

(a) Explain why a horizontally rotating wire carrying a mass at its free end can break at a rotation speed at which an equivalent static tensile load would be perfectly safe.

(b) A copper wire of length 1.0m and diameter 0.80mm has a mass of 0.50kg attached to its free end. The other end is held fixed and the wire is rotated in a horizontal circle. The breaking stress of the copper wire is $2.5 \times 10^{8}\text{Pa}$. Ignoring the mass

of the wire itself, calculate (i) the maximum tension the wire can withstand just before breaking, (ii) the maximum angular velocity that the wire can survive, and (iii) the corresponding linear speed of the mass at this maximum.

Solution

(a) A rotating wire is subjected to repeated dynamic stressing accompanied by small bending and vibrational stresses that a pure static test does not reproduce. Elastic fatigue accumulates microscopic cracks at surface flaws and lowers the effective breaking stress well below its static value. The wire therefore fails sooner under rotation than under an equivalent static load.

(b) (i) Cross-sectional area of the wire:

$$A = \frac{\pi d^2}{4} = \frac{\pi(0.80 \times 10^{-3}\text{m})^2}{4} = 5.03 \times 10^{-7}\text{m}^2$$

Maximum tension:

$$T_{max} = \sigma_{break} \cdot A = (2.5 \times 10^8\text{Pa})(5.03 \times 10^{-7}\text{m}^2) = 125.7\text{N}$$

(ii) The tension provides the centripetal force $T = m\omega^2 r$, with $r = L = 1.0\text{m}$:

$$\omega_{max} = \sqrt{\frac{T_{max}}{mr}} = \sqrt{\frac{125.7\text{N}}{(0.50\text{kg})(1.0\text{m})}} = 15.9\text{rad/s}$$

(iii) Linear speed of the mass at this angular velocity:

$$v = \omega r = (15.9\text{rad/s})(1.0\text{m}) = 15.9\text{m/s}$$

Example 35

(a) A long rubber band and a short steel wire of the same cross-sectional area are each stretched until they store the same amount of elastic potential energy per unit volume. Which of the two is strained more, and state the physical reason.

(b) A steel wire of length 1.50m, cross-sectional area $2.0 \times 10^{-7}\text{m}^2$, and Young's modulus $2.0 \times 10^{11}\text{Pa}$ is stretched by an axial force until its strain reaches 1.0×10^{-3}. (i) Calculate the tensile stress in the wire. (ii) Calculate the extension of the wire. (iii) Calculate the elastic energy per unit volume stored in the wire.

Solution

(a) The rubber band is strained more. The elastic energy per unit volume is $u = \frac{1}{2}E\varepsilon^2$. For the two materials to store the same energy per unit volume, the product $E\varepsilon^2$ must be the same. Since steel is much stiffer, the rubber's Young's modulus is many orders of magnitude smaller than steel's; so the rubber's strain must be many times larger to compensate.

(b) (i) Tensile stress in the wire:

$$\sigma = E\varepsilon = (2.0 \times 10^{11}\text{Pa})(1.0 \times 10^{-3}) = 2.0 \times 10^8\text{Pa}$$

(ii) Extension of the wire:

$$\Delta L = \varepsilon L = (1.0 \times 10^{-3})(1.50\text{m}) = 1.50 \times 10^{-3}\text{m} = 1.50\text{mm}$$

(iii) Elastic energy per unit volume:

$$u = \frac{1}{2}\sigma\varepsilon = \frac{1}{2}(2.0 \times 10^8\text{Pa})(1.0 \times 10^{-3}) = 1.0 \times 10^5\text{J/m}^3$$

Example 36

(a) Explain why a long vertical wire can break under its own weight even when no external load is attached to its lower end.

(b) A steel wire of density 7800kg/m^3 and breaking stress $8.0 \times 10^8\text{Pa}$ is suspended vertically from a rigid support. Take $g = 9.8\text{m/s}^2$. (i) Derive an expression for the maximum length of such a wire that can hang vertically without breaking at its point of suspension. (ii) Calculate the numerical value of this maximum length.

Solution

(a) The point of suspension at the top of the wire has to carry the entire weight of the wire hanging below it. As the length of the wire increases, the weight below the top increases in proportion, and the tensile stress at the top increases with it. When the stress at the top reaches the breaking stress of the material, the wire snaps at that point. The rest of the wire, which carried less weight, remains intact.

(b) (i) Consider a wire of length L, cross-sectional area A, and density ρ. The total weight of the wire is:

$$W = \rho A L g$$

All weight acts on the top cross-section. The stress at the top is therefore:

$$\sigma_{top} = \frac{W}{A} = \rho L g$$

Setting σ_{top} equal to the breaking stress σ_{break} and solving for the maximum length:

$$\boldsymbol{L_{max} = \frac{\sigma_{break}}{\rho g}}$$

(ii) Substituting the values:

$$L_{max} = \frac{8.0 \times 10^8 \text{Pa}}{(7800 \text{kg/m}^3)(9.8 \text{m/s}^2)} = 1.05 \times 10^4 \text{m}$$

Example 37

(a) Explain why a fluid cannot support a static shear stress in equilibrium, while a solid of the same volume can.

(b) A cube of steel of side 5.0cm is subjected to a shearing force of 2.5×10^5N applied tangentially along one of its faces while the opposite face is held fixed. The shear modulus of steel is 8.0×10^{10}Pa. Calculate (i) the shear stress on the face, (ii) the shear strain in the cube, and (iii) the lateral displacement of the top face relative to the bottom face.

Solution

(a) In a fluid, the layers have no binding force resisting tangential displacement relative to one another. When a shear stress is applied, the layers slide past each other continuously, and the fluid flows rather than settling into equilibrium. In a solid, the intermolecular forces exert restoring forces against tangential displacement, and the solid reaches a static equilibrium at a finite shear strain determined by the shear modulus.

(b) (i) Area of the sheared face:

$$A = (5.0 \times 10^{-2} \text{m})^2 = 2.5 \times 10^{-3} \text{m}^2$$

Shear stress:

$$\tau = \frac{F}{A} = \frac{2.5 \times 10^5 \text{N}}{2.5 \times 10^{-3} \text{m}^2} = 1.0 \times 10^8 \text{Pa}$$

(ii) Shear strain:

$$\gamma = \frac{\tau}{G} = \frac{1.0 \times 10^8 \text{Pa}}{8.0 \times 10^{10} \text{Pa}} = 1.25 \times 10^{-3}$$

(iii) Lateral displacement of the top face:

$$x = \gamma h = (1.25 \times 10^{-3})(5.0 \times 10^{-2} \text{m}) = 6.25 \times 10^{-5} \text{m} = 0.0625 \text{mm}$$

Example 38

(a) State the physical meaning of the compressibility of a fluid, and explain why gases have much higher compressibilities than liquids at ordinary pressures.

(b) A volume of 1.025m^3 of sea water is taken from the surface to a depth where the pressure is 5.0×10^7Pa greater than at the surface. The bulk modulus of sea water is 2.2×10^9Pa, and the density of sea water at the surface is 1025kg/m^3. Calculate the density of the sea water at this greater depth.

Solution

(a) Compressibility is the fractional change in volume of the fluid per unit increase in applied pressure; it is the reciprocal of the bulk modulus. Gases have much higher compressibilities than liquids because their molecules are far apart and there is large empty space between them that can be squeezed out by an applied pressure. In liquids the molecules are almost in contact and their mutual repulsion makes further compression extremely difficult, so liquids are nearly incompressible.

(b) *Calculating the fractional change in volume:*

$$\frac{\Delta V}{V} = -\frac{\Delta P}{B} = -\frac{5.0 \times 10^7 \text{Pa}}{2.2 \times 10^9 \text{Pa}} = -2.27 \times 10^{-2}$$

Calculating the actual change in volume:

$$\Delta V = -2.27 \times 10^{-2} \times 1.025 \text{m}^3 = -2.33 \times 10^{-2} \text{m}^3$$

The mass of the sea water does not change with compression. Mass at the surface:

$$m = \rho_0 V_0 = (1025 \text{kg/m}^3)(1.025 \text{m}^3) = 1050.6 \text{kg}$$

Calculating new volume at the depth:

$$V' = 1.025 \text{m}^3 - 2.33 \times 10^{-2} \text{m}^3 = 1.0017 \text{m}^3$$

New density at the depth:

$$\rho' = \frac{m}{V'} = \frac{1050.6 \text{kg}}{1.0017 \text{m}^3} = 1049 \text{kg/m}^3$$

Alternative solution

For constant mass, $\rho \propto \frac{1}{V}$. So fractional decrease in volume = fractional **increase** in density

$$\rho' = \rho_0 + \Delta\rho = 1025 \text{kg/m}^3 + 2.27 \times 10^{-2} \times 1025 \text{kg/m}^3 = 1048 \text{kg/m}^3$$

Example 39

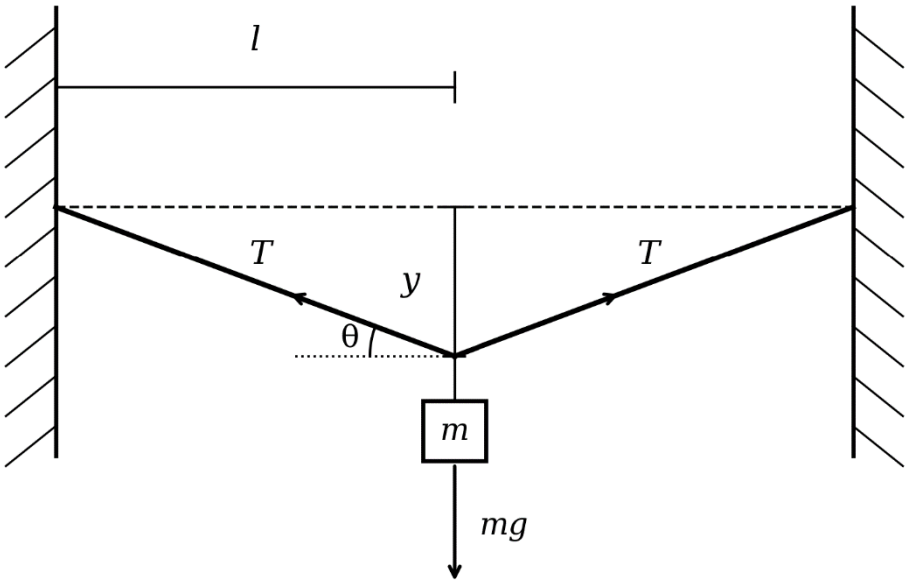
(a) Explain why a horizontal wire supporting a small load at its midpoint cannot remain perfectly horizontal, however short the wire or however small the load.

(b) A mass of 1.0kg is suspended from the midpoint of a steel wire of total length 2.0m stretched horizontally between two rigid supports. The wire has cross-sectional area $2.0 \times 10^{-6}\text{m}^2$ and Young's modulus $2.0 \times 10^{11}\text{Pa}$. Take $g = 9.8\text{m/s}^2$. (i) Derive, using the small-angle approximation, an expression relating the midpoint depression y to the mass m, the half-length l of the wire, the cross-sectional area A, and the Young's modulus E. (ii) Calculate the numerical value of the midpoint depression.

Solution

(a) A perfectly horizontal wire can only exert tension along its own length, which is horizontal. A horizontal tension has no vertical component, and cannot support the downward weight of the suspended load. The wire must therefore tilt downward slightly at its midpoint so that its tension acquires a small vertical component, which is just large enough to balance the weight of the load. The midpoint must sag.

(b) (i) Consider the following diagram:



(b) (i) Two methods of derivation are presented below; you are invited to study both and decide which one feels more natural.

Method 1: Pythagoras and the binomial approximation

Finding the new length of each half

After the wire sags, each half runs from a fixed support down to the midpoint, forming the hypotenuse of a right-angled triangle whose horizontal side is l and whose vertical side is y. By Pythagoras' theorem, the new length of each half is:

$$L = \sqrt{l^2 + y^2}$$

Applying the small-sag approximation by factoring out l

Since y is much smaller than l, we factor l out of the square root to bring the small quantity into a recognisable form:

$$L = \sqrt{l^2\left(1 + \frac{y^2}{l^2}\right)} = l\sqrt{1 + \frac{y^2}{l^2}}$$

For any small quantity x, the binomial approximation gives $\sqrt{1+x} \approx 1 + \frac{1}{2}x$. Applying this with $x = \frac{y^2}{l^2}$ (which is very small because $y \ll l$):

$$L \approx l\left(1 + \frac{1}{2} \times \frac{y^2}{l^2}\right) = l + \frac{y^2}{2l}$$

Finding the extension of each half

The extension of each half is its new length minus its original length l:

$$\Delta l = L - l = \left(l + \frac{y^2}{2l}\right) - l = \frac{y^2}{2l}$$

Finding the strain in each half

$$\varepsilon = \frac{\Delta l}{l} = \frac{\frac{y^2}{2l}}{l} = \frac{y^2}{2l^2}$$

Finding the tension in each half from Young's modulus

The stress in each half is $\sigma = E\varepsilon = \frac{Ey^2}{2l^2}$, and the tension is the stress times the cross-sectional area:

$$T = \sigma A = \frac{EAy^2}{2l^2}$$

Applying equilibrium at the midpoint

The midpoint is in equilibrium under three forces: the two tensions pointing along the wire from the midpoint outward toward the supports, and the weight mg pointing downward. The two horizontal components of tension cancel each other by symmetry. The two vertical components must together support the weight. Each vertical component is $T\sin\theta$, where θ is the small angle each half of the wire makes with the horizontal. From the geometry, $\sin\theta \approx \tan\theta = \frac{y}{l}$ for small θ, so:

$$2T \times \frac{y}{l} = mg$$

Substituting T from $T = \sigma A = \frac{EAy^2}{2l^2}$ and solve for y

$$2 \times \frac{EAy^2}{2l^2} \times \frac{y}{l} = mg$$

$$\frac{EAy^3}{l^3} = mg$$

$$y^3 = \frac{mgl^3}{EA}$$

$$\mathbf{y = l\sqrt[3]{\frac{mg}{EA}}}$$

Method 2: Direct small-angle method

Defining the small angle and write its trigonometric ratios

Let θ be the small angle that each half of the wire makes with the horizontal after sagging. From the geometry:

$$\tan\theta = \frac{y}{l}$$

For a small angle θ, the standard small-angle approximations are:

$$\sin\theta \approx \tan\theta \approx \theta, \qquad \cos\theta \approx 1 - \frac{\theta^2}{2}$$

In particular, $\sin\theta \approx \frac{y}{l}$

Finding the new length of each half from the cosine

The horizontal half-length l is the projection of the new (stretched) half-length L onto the horizontal direction:

$$l = L\cos\theta$$

Solving for L:

$$L = \frac{l}{\cos\theta}$$

For a small angle, $\frac{1}{\cos\theta} \approx 1 + \frac{\theta^2}{2}$, so:

$$L \approx l\left(1 + \frac{\theta^2}{2}\right) = l + \frac{l\theta^2}{2}$$

Finding the extension and the strain

$$\Delta l = L - l = \frac{l\theta^2}{2}$$

$$\varepsilon = \frac{\Delta l}{l} = \frac{\theta^2}{2}$$

Substituting $\theta \approx \frac{y}{l}$:

$$\varepsilon = \frac{1}{2}\left(\frac{y}{l}\right)^2 = \frac{y^2}{2l^2}$$

Finding the tension in each half from Young's modulus

$$T = \sigma A = E\varepsilon A = \frac{EAy^2}{2l^2}$$

Applying equilibrium at the midpoint

The two vertical components of the tensions must support the weight mg:

$$2T\sin\theta = mg$$

Substituting $\sin\theta \approx y/l$ from Step 1:

$$2T \cdot \frac{y}{l} = mg$$

Substituting T and solve for y

$$2 \cdot \frac{EAy^2}{2l^2} \cdot \frac{y}{l} = mg$$

$$\frac{EAy^3}{l^3} = mg$$

$$\mathbf{y = l\sqrt[3]{\frac{mg}{EA}}}$$

The two methods reach the same expression by different routes. Method 1 stays in pure algebra and applies the small-quantity approximation to a square root. Method 2 stays in pure trigonometry and applies the small-angle approximations. Both are equally valid; whichever feels more natural is the one to use in the examination room.

(ii) Substituting:

$$y = (1.0\text{m})\sqrt[3]{\frac{(1.0\text{kg})(9.8\text{m/s}^2)}{(2.0\times10^{11}\text{Pa})(2.0\times10^{-6}\text{m}^2)}} = 2.90\times10^{-2}\text{m} = 2.9\text{cm}$$

The midpoint of the wire therefore sags by approximately 2.9cm under the suspended kilogram mass.

Example 40

(a) Two wires of different materials are clamped between the same pair of rigid walls so that they lie side by side, each running the full distance between the walls. Explain what happens to each wire as the temperature of the whole system is lowered below the reference temperature at which the wires were originally fitted.

(b) A steel wire and a copper wire, each of length 50.0cm at 20°C, are clamped between the same two rigid walls. The system is then cooled to 0°C. Given $E_s = 2.0\times10^{11}\text{Pa}$, $E_c = 1.2\times10^{11}\text{Pa}$, $\alpha_s = 12\times10^{-6}\text{K}^{-1}$, and $\alpha_c = 17\times10^{-6}\text{K}^{-1}$, calculate (i) the tensile thermal stress developed in each wire, and (ii) state which wire experiences the greater thermal stress and identify the physical reason.

Solution

(a) When the temperature is lowered, each wire tries to contract by an amount proportional to its length, its coefficient of linear expansion, and the temperature drop. The walls hold the ends of both wires fixed, so neither can shorten. Each wire therefore develops a tensile stress large enough to stretch the material elastically by exactly the amount it wanted to contract thermally. Because the two wires have different Young's moduli and different coefficients of expansion, they develop different tensile stresses even though they experience the same temperature drop.

(b) (i) The magnitude of the temperature change is $|\Delta T| = 20\text{K}$. Thermal stress in each wire is $\sigma = E\alpha|\Delta T|$.
For the steel wire:

$$\sigma_s = E_s\alpha_s|\Delta T| = (2.0\times10^{11}\text{Pa})(12\times10^{-6}\text{K}^{-1})(20\text{K}) = 4.8\times10^{7}\text{Pa}$$

For the copper wire:

$$\sigma_c = E_c\alpha_c|\Delta T| = (1.2\times10^{11}\text{Pa})(17\times10^{-6}\text{K}^{-1})(20\text{K}) = 4.08\times10^{7}\text{Pa}$$

(ii) The steel wire experiences the greater thermal stress ($4.8\times10^{7}\text{Pa}$ compared with $4.08\times10^{7}\text{Pa}$ for copper). Although copper has a larger coefficient of expansion, steel has a much larger Young's modulus, and the product $E\alpha$ for steel exceeds that of copper, so steel is the stiffer partner in the pair.

Example 41

(a) Explain why, in the design of a long railway line, the length of individual continuous rail sections is limited not by the breaking stress of the steel but by the maximum expected temperature change at the location.

(b) A steel rod is clamped rigidly between two walls at 15°C. The maximum compressive stress that the rod can safely withstand before it is judged to buckle is $1.5\times10^{8}\text{Pa}$. Given $E = 2.0\times10^{11}\text{Pa}$ and $\alpha = 12\times10^{-6}\text{K}^{-1}$, (i) derive an expression for the maximum temperature rise above 15°C at which the rod remains safe, (ii) calculate this maximum temperature rise numerically, and (iii) state the highest safe ambient temperature for the rod.

Solution

(a) A clamped rail develops a thermal stress $\sigma = E\alpha\Delta T$ that is independent of its length. However, a long slender rail does not fail by direct compression at the full crushing strength of steel: it fails by sideways buckling at stresses much smaller than the crushing strength. The longer the rail, the lower the buckling stress. Expansion gaps must therefore be sized so that they always close before buckling begins, and this requirement limits the length of any continuous rail section based on the maximum temperature change the rail can experience between the coldest night and the hottest afternoon.

(b) (i) Setting $\sigma = E\alpha\Delta T$ equal to the maximum safe compressive stress and solving for ΔT:

$$\Delta T_{max} = \frac{\sigma_{max}}{E\alpha}$$

(ii) Substituting the values:

$$\Delta T_{max} = \frac{1.5 \times 10^8 \text{Pa}}{(2.0 \times 10^{11} \text{Pa})(12 \times 10^{-6} \text{K}^{-1})} = 62.5\text{K}$$

(iii) The highest safe ambient temperature is:

$$T_{max} = 15°\text{C} + 62.5°\text{C} = 77.5°\text{C}$$

Example 42

(a) Starting from the definition of work done by a variable force, derive the expression $U = \frac{1}{2}Fx$ for the elastic potential energy stored in a wire that obeys Hooke's law, where F is the final applied force and x is the corresponding extension.

(b) A brass wire of length 0.50m, cross-sectional area $1.5 \times 10^{-6}\text{m}^2$, and Young's modulus $9.0 \times 10^{10}\text{Pa}$ is stretched by a force of 80N. Calculate (i) the extension of the wire, and (ii) the elastic potential energy stored in the wire.

Solution

(a) For a wire obeying Hooke's law, the force required to produce an extension x' is $F' = kx'$, where k is the spring constant of the wire. The work done in increasing the extension from x' to $x' + dx'$ against this variable force is $dW = F'\,dx' = kx'\,dx'$. Integrating from the unstretched state $x' = 0$ to the final extension $x' = x$:

$$W = \int_0^x k\,x'\,dx' = \frac{1}{2}kx^2$$

The spring constant can be eliminated using $k = \frac{F}{x}$ at the final extension, giving $W = \frac{1}{2}\left(\frac{F}{x}\right)x^2 = \frac{1}{2}Fx$.

Since all the work done against the elastic force is stored as elastic potential energy, $U = W$, therefore:

$$\mathbf{U = \frac{1}{2}Fx}$$

(b) (i) Extension of the wire:

$$\Delta L = \frac{FL}{AE} = \frac{(80\text{N})(0.50\text{m})}{(1.5 \times 10^{-6}\text{m}^2)(9.0 \times 10^{10}\text{Pa})} = 2.96 \times 10^{-4}\text{m}$$

(ii) Elastic potential energy stored:

$$U = \frac{1}{2}F\Delta L = \frac{1}{2}(80\text{N})(2.96 \times 10^{-4}\text{m}) = 1.19 \times 10^{-2}\text{J}$$

Example 43

(a) Explain what is meant by stress concentration in an engineering component, give one example of its consequences in practice, and state the design modification commonly used to reduce it.

(b) A steel wire of uniform diameter 1.0mm is used to suspend a mass vertically. The ultimate tensile strength of the steel is $8.0 \times 10^8\text{Pa}$ and the required safety factor is 5. Take $g = 9.8\text{m/s}^2$. Calculate (i) the maximum working stress in the wire, (ii) the maximum mass that the wire may safely suspend, and (iii) the new maximum safe mass if a surface defect reduces the effective diameter of the wire to 0.9mm at one point, assuming the stress is concentrated at the narrowest cross-section.

Solution

(a) Stress concentration is the localised amplification of stress at a geometric discontinuity in a component, such as a sharp corner, a notch, a hole, or a sudden change in cross-section. The stress at the discontinuity can be several times larger than the nominal stress elsewhere, so the component can fail at that point while the average stress is still well below the breaking stress. The classic example is the de Havilland Comet aircraft of the 1950s, three of which broke up in flight because fatigue cracks developed at the sharp corners of their rectangular windows. The design modification is to round off every sharp corner and to smooth any sudden change in cross-section, so that stress varies gradually rather than jumping at an edge.

(b) (i) Maximum working stress:

$$\sigma_{working} = \frac{\sigma_{UTS}}{n} = \frac{8.0 \times 10^8\text{Pa}}{5} = 1.6 \times 10^8\text{Pa}$$

(ii) Original cross-sectional area:

$$A = \frac{\pi d^2}{4} = \frac{\pi(1.0 \times 10^{-3}\text{m})^2}{4} = 7.85 \times 10^{-7}\text{m}^2$$

Maximum force the wire may carry:

$$F_{max} = \sigma_{working} \cdot A = (1.6 \times 10^8\text{Pa})(7.85 \times 10^{-7}\text{m}^2) = 125.7\text{N}$$

Maximum mass:

$$m_{max} = \frac{F_{max}}{g} = \frac{125.7\text{N}}{9.8\text{m/s}^2} = 12.8\text{kg}$$

(iii) Reduced area at the defect:

$$A' = \frac{\pi(0.9 \times 10^{-3}\text{m})^2}{4} = 6.36 \times 10^{-7}\text{m}^2$$

New maximum force and mass:

$$F'_{max} = (1.6 \times 10^8 Pa)(6.36 \times 10^{-7} m^2) = 101.8N$$

$$m'_{max} = \frac{101.8N}{9.8 m/s^2} = 10.4kg$$

Example 44

(a) Two vertical rods of different materials support a rigid horizontal beam side by side. Explain why, when a load is suspended from the beam, the two rods carry different shares of the load even though the beam is rigid and horizontal.

(b) A copper rod of cross-sectional area $2.0 \times 10^{-6} m^2$ and a steel rod of cross-sectional area $1.0 \times 10^{-6} m^2$, each of length 1.5m, support a rigid horizontal beam in parallel. A total load of 1200N is shared between the two rods. $E_c = 1.2 \times 10^{11} Pa$, $E_s = 2.0 \times 10^{11} Pa$. Assume the beam remains horizontal. Calculate (i) the force supported by each rod, and (ii) the common extension of the two rods.

Solution

(a) Because the beam is rigid and horizontal, both rods must extend by the same amount. The force required to produce a given extension in a rod is proportional to its Young's modulus and its cross-sectional area (for the same length), so the rod with the larger product EA carries a proportionally larger share of the load. The two rods therefore carry different forces even though they experience the same extension, and the distribution of the load is determined by the relative stiffness of the rods.

(b) (i) Let F_c and F_s be the forces in the copper and steel rods. Equal extension requires:

$$\frac{F_c L}{A_c E_c} = \frac{F_s L}{A_s E_s}$$

The common length L cancels, giving:

$$\frac{F_s}{F_c} = \frac{A_s E_s}{A_c E_c} = \frac{(1.0 \times 10^{-6} m^2)(2.0 \times 10^{11} Pa)}{(2.0 \times 10^{-6} m^2)(1.2 \times 10^{11} Pa)} = 0.833$$

Combined with $F_c + F_s = 1200N$:

$$F_c(1 + 0.833) = 1200N \quad \Rightarrow \quad F_c = 654.6N, \quad F_s = 545.4N$$

(ii) Common extension, calculated from the copper rod:

$$\Delta L = \frac{F_c L}{A_c E_c} = \frac{(654.6N)(1.5m)}{(2.0 \times 10^{-6} m^2)(1.2 \times 10^{11} Pa)} = 4.09 \times 10^{-3} m = 4.09mm$$

You may check with the steel rod, and will give the same value, confirming consistency.

Example 45

(a) Explain the physical principle underlying the shrink-fitting process, in which a metal ring is joined to a metal shaft without bolts, welds, or keys.

(b) A brass ring of internal diameter 25.00mm at 20°C is to be fitted over a steel shaft of external diameter 25.05mm at the same temperature. The coefficient of linear expansion of brass is $19 \times 10^{-6} K^{-1}$. (i) Calculate the minimum temperature to which the brass ring must be heated so that its internal diameter just equals the diameter of the shaft. (ii) Describe what happens as the ring then cools back to 20°C while in position on the shaft.

Solution

(a) A metal ring expands when heated and contracts when cooled. By heating the ring until its internal diameter is slightly larger than the external diameter of the shaft, the ring can be slipped onto the shaft without force. As the ring cools back to ambient temperature, it tries to contract back to its original smaller internal diameter, but the shaft prevents the inner surface from shrinking. The ring therefore develops a large tensile hoop stress and grips the shaft with a radial compressive force that is strong enough to transmit torques without any keys, bolts, or welds.

(b) (i) The internal diameter of the ring must increase from 25.00mm to 25.05mm, an increase of 0.05mm. Using the linear expansion formula for the diameter (the internal diameter expands as a solid piece of brass would):

$$\Delta d = d_0 \alpha \Delta T \quad \Rightarrow \quad \Delta T = \frac{\Delta d}{d_0 \alpha}$$

Substituting the numerical values:

$$\Delta T = \frac{0.05 \times 10^{-3} m}{(25.00 \times 10^{-3} m)(19 \times 10^{-6} K^{-1})} = 105.3K$$

The ring must therefore be heated to:

$$T = 20°C + 105.3°C \approx 125°C$$

(ii) As the ring cools back to 20°C while seated on the shaft, it tries to return to its original smaller diameter, but the shaft prevents the inner surface of the ring from contracting below 25.05mm. The ring therefore remains stretched by 0.05mm in

diameter relative to its natural size, and develops a tensile hoop stress that presses its inner surface radially inward against the shaft with a large compressive force. This force holds the two components together permanently.

Example 46

(a) A steel wire is simultaneously loaded with a tensile force and heated. Explain which contribution to the total strain obeys Hooke's law and which does not, and state whether the two strains add or subtract when the wire is heated.

(b) A steel wire of length 2.0m and cross-sectional area $1.0 \times 10^{-6}\text{m}^2$ is stretched by a load of 100N at $20°\text{C}$, and the system is then heated to $50°\text{C}$ while the load is maintained. $E_{steel} = 2.0 \times 10^{11}\text{Pa}$, $\alpha_{steel} = 12 \times 10^{-6}\text{K}^{-1}$. Calculate (i) the elastic strain produced by the load alone at $20°\text{C}$, (ii) the thermal strain produced by the temperature rise, and (iii) the total extension of the wire at $50°\text{C}$.

Solution

(a) The elastic strain, which is produced by the applied tensile force, obeys Hooke's law: it is proportional to the stress through Young's modulus, $\varepsilon_{elastic} = \sigma/E$. The thermal strain, produced by the temperature change, does not obey Hooke's law: it is given by $\varepsilon_{thermal} = \alpha\Delta T$ and is independent of the applied stress. When a wire is both stretched and heated, both effects cause the wire to lengthen, so the two strains add together.

Understand! If the wire were cooled instead of heated, the thermal strain would be negative and would subtract from the elastic strain.

(b) (i) Tensile stress and elastic strain from the load:

$$\sigma = \frac{F}{A} = \frac{100\text{N}}{1.0 \times 10^{-6}\text{m}^2} = 1.0 \times 10^{8}\text{Pa}$$

$$\varepsilon_{elastic} = \frac{\sigma}{E} = \frac{1.0 \times 10^{8}\text{Pa}}{2.0 \times 10^{11}\text{Pa}} = 5.0 \times 10^{-4}$$

(ii) Thermal strain from the temperature rise of $\Delta T = 30\text{K}$:

$$\varepsilon_{thermal} = \alpha\Delta T = \left(12 \times 10^{-6}\text{K}^{-1}\right)(30\text{K}) = 3.6 \times 10^{-4}$$

(iii) Total strain and total extension at $50°\text{C}$:

$$\varepsilon_{total} = \varepsilon_{elastic} + \varepsilon_{thermal} = 5.0 \times 10^{-4} + 3.6 \times 10^{-4} = 8.6 \times 10^{-4}$$

$$\Delta L = \varepsilon_{total} \cdot L = (8.6 \times 10^{-4})(2.0\text{m}) = 1.72 \times 10^{-3}\text{m} = 1.72\text{mm}$$

Forty-six worked examples now lie behind you, and thirty-three fresh questions wait in the Digging Deeper Exercise that follows.

DIGGING DEEPER EXERCISE 1

EXERCISE 1A: BINDER QUESTIONS

Question 1

A wire that obeys Hooke's law extends by ΔL when a force F is applied. Explain in physical terms why doubling the force doubles the extension only as long as the stress remains below a certain limit, and identify what that limit is called.

Question 2

A long thin wire and a short thick wire are made from the same material and stretched by the same force. Explain which wire extends more and identify the two geometric factors that determine the answer.

Question 3

Distinguish between Young's modulus, the shear modulus, and the bulk modulus of a material. For each modulus, identify the type of deformation it describes and the geometric situation in which it is the relevant material property.

Question 4

Explain why a fluid has a finite bulk modulus but no shear modulus, and use your explanation to identify the essential physical difference between a solid and a fluid.

Question 5

A rubber band and a steel wire of the same dimensions are each stretched by the same small force. Explain which one extends more, and explain which one stores more elastic potential energy.

Question 6

Distinguish between the elastic potential energy stored in a stretched wire and the energy stored per unit volume of the wire. Explain why the energy per unit volume is the more useful quantity when comparing two materials.

Question 7

Explain what is meant by the elastic fatigue of a material, and describe how it differs from ordinary overloading. Give one example of an everyday object whose useful life is limited by elastic fatigue rather than by overloading.

Question 8

Sketch a typical stress-strain curve for a ductile metal and label on it the proportional limit, the elastic limit, the yield point, the ultimate tensile strength, and the fracture point. Explain in one sentence what is happening to the metal between each pair of adjacent labelled points.

Question 9

A bimetallic strip is formed by bonding a strip of brass to a strip of steel of equal length. Explain why the strip bends when heated, identify which metal lies on the outer side of the curve and explain why, and describe one practical device whose operation depends on this bending.

Question 10

Explain the working principle of Searle's two-wire apparatus for determining the Young's modulus of a wire, and identify two advantages of using a two-wire arrangement over a single-wire experiment.

EXERCISE 1B: REAL QUESTIONS

Question 11

Mama Amina buys a new spring balance from a hardware shop in Kariakoo and uses it daily to weigh sugar in her shop. After two years, she notices that the balance reads 50g even when nothing is hanging from it, although the spring still works smoothly and visibly returns to its same position whenever she removes a load. Explain in physical terms what has happened to the spring, and explain why simply applying a correction of -50g to every reading is not a complete solution to the problem.

Question 12

Explain why expansion joints are built into long suspension bridges in coastal cities.

Question 13

A car tyre is much softer than a metal wheel rim, and engineers deliberately design tyres that way rather than making them as stiff as the rim. Explain in physical terms what role the elastic compliance of the tyre plays in the everyday operation of a car, and explain why a car with completely rigid wheels would be intolerable to drive even on a smooth road.

Question 14

Explain why an old riveted steel bridge creaks under load while a similar welded bridge does not.

Question 15

Explain why steel reinforcing bars in a concrete beam are placed near the bottom of the beam rather than near the top.

Question 16

Explain why a motorcycle helmet must contain a layer of crushable foam in addition to its hard outer shell.

Question 17

Mr. Akilikubwa has just sent Kipanga to the staff room to fetch a textbook. Kipanga returns a few minutes later, looking pleased with himself.

Kipanga: *"Sir, somebody has broken your window in the staff room! There is a long crack running from the corner all the way to the middle of the glass. It must have been one of those Form Three boys."*

Kipute: *"Are you sure, Kipanga? I was in the staff room yesterday and the crack was already there. Did you look carefully at where it started?"*

Kipanga: *"It started from the corner. So what?"*

Kipute: *"And tell me, when the carpenter installed that window last term, did he put rubber spacers around the edges of the glass before fixing the frame, or did he press the glass straight against the wood?"*

Kipanga: *"I do not remember. Why does it matter?"*

Explain why the crack started at the corner of the window, and identify the role that the missing rubber spacers played in producing it.

Question 18

Kipanga goes fishing at Kunduchi beach with a stiff nylon line. He hooks a large fish, and as he pulls the rod sharply upward to set the hook, the line snaps even though the breaking strength of the line is well above the weight of the fish. Explain why a sudden jerk on a stiff line is more likely to break it than a steady pull of the same average force.

Question 19

Mr. Akilikubwa is sitting in the staff room at break time, marking a stack of test papers, when Kipute comes in looking troubled.

Kipute: *"Sir, my grandmother fell at home last week and broke her hip. The doctor at the hospital said it is very common in elderly people. But I have been reading about elasticity in the textbook, and it says that bone has a high Young's modulus, comparable to that of concrete. If bone is so stiff and so strong, how can it break from such a small fall? And another thing Sir, why was it her hip and not her wrist or her arm? She was holding a teacup when she fell, but the teacup did not break and her arm was not injured."*

Mr. Akilikubwa puts down his red pen and looks at Kipute thoughtfully for a long moment before answering.

Explain why an elderly person's bone can fracture under a load that a young person's bone of the same Young's modulus would survive.

EXERCISE 1C: HOT QUESTIONS

Question 20

A long steel cable used in a colliery winding gear has cross-sectional area $4.0 \times 10^{-4}\text{m}^2$ and supports a cage of mass 1500kg that descends at 3.0m/s and is brought to rest in 0.40s by the action of the cable. Taking $g = 9.8\text{m/s}^2$, calculate the maximum stress developed in the cable during the deceleration.

Question 21

A brass ring with internal diameter 30.00mm at $20°\text{C}$ is to be shrink-fitted onto a steel shaft of external diameter 30.06mm at the same temperature. Given the coefficient of linear expansion of brass as $19 \times 10^{-6}\text{K}^{-1}$, calculate the temperature to which the brass ring must be heated so that it just slips over the shaft.

Question 22

A vertical steel wire just breaks under its own weight when it has length $1.05 \times 10^{4}\text{m}$. Taking $g = 9.8\text{m/s}^2$ and the density of steel as 7800kg/m^3, calculate the breaking stress of the steel.

Question 23

A copper wire of length 1.0m and diameter 0.80mm is rotated in a horizontal circle about one of its ends. A mass of 0.40kg is attached to the free end of the wire, and the breaking stress of the copper is $2.5 \times 10^{8}\text{Pa}$. Calculate the maximum linear speed of the mass at the moment the wire breaks.

Question 24

Two vertical wires hang side by side from the same rigid support. Wire A is steel with $E_s = 2.0 \times 10^{11}$Pa, length 1.0m and cross-sectional area 1.0×10^{-6}m^2. Wire B is brass with $E_b = 9.0 \times 10^{10}$Pa, length 1.5m and cross-sectional area 2.0×10^{-6}m^2. The same mass is hung from each wire in turn. Calculate the ratio of the elastic potential energy stored in wire A to that stored in wire B.

Question 25

A solid steel cube of side 8.0cm at 25°C is held inside a perfectly rigid jig that prevents it from expanding in any direction. The cube is heated to 95°C. Given the bulk modulus of steel as 1.6×10^{11}Pa and $\alpha = 12 \times 10^{-6}$K^{-1}, calculate the total compressive force exerted by the jig on each face of the cube.

Question 26

A brass ring is heated from 20°C to a maximum safe temperature of 200°C in order to be shrink-fitted onto a steel shaft. The coefficient of linear expansion of brass is 19×10^{-6}K^{-1}. If the internal diameter of the ring at 20°C is 25.00mm, calculate the maximum diameter of steel shaft that the heated ring can be slipped over.

Question 27

A steel rod is clamped rigidly between two walls at 20°C. The Young's modulus of steel is 2.0×10^{11}Pa, its coefficient of linear expansion is 12×10^{-6}K^{-1}, and its ultimate compressive strength is 5.0×10^{8}Pa. Calculate the maximum temperature to which the rod can be heated before it crushes.

Question 28

A passenger lift uses 6 identical steel cables in parallel, each of diameter 5.0mm. The yield stress of the steel is 2.5×10^{8}Pa and the safety factor required by the building code is 5. Taking $g = 9.8$m/s^2, calculate the maximum live load (in kilograms) that the lift may safely carry.

Question 29

A horizontal steel wire of total length 2.0m and cross-sectional area 1.5×10^{-6}m^2 is stretched between two rigid supports, and a mass is suspended from its midpoint. The midpoint is observed to sag by 5.0cm below its original horizontal position. Given $E = 2.0 \times 10^{11}$Pa and $g = 9.8$m/s^2, calculate the mass of the suspended load.

Question 30

A copper wire of length 80cm and diameter 0.50mm is attached to a 0.20kg mass and used as a conical pendulum, with the wire making a steady angle of 30° with the vertical. Given that $E_c = 1.2 \times 10^{11}$Pa and $g = 9.8$m/s^2, calculate the extension of the wire as the pendulum rotates steadily.

Question 31

A diver descends below the surface of a lake to a depth at which the density of the water has increased to 1003.0kg/m^3 from its surface value of 1000kg/m^3. The bulk modulus of the water is 2.2×10^{9}Pa, and $g = 9.8$m/s^2. Calculate the depth of the diver below the surface.

Question 32

A 2.0kg steel ball is attached to the lower end of a vertical steel wire of length 3.0m and cross-sectional area 1.0×10^{-6}m^2, the upper end of the wire being fixed to a rigid support. The ball is held so that the wire is just taut and at its natural length, and is then released so that it falls and is caught by the wire. Treating the ball-and-wire system as conservative and taking $E = 2.0 \times 10^{11}$Pa and $g = 9.8$m/s^2, calculate the maximum extension of the wire during the subsequent motion.

Question 33

A rubber block of dimensions 2.0cm × 2.0cm × 5.0cm is glued by one of its 2.0cm × 2.0cm faces to a fixed wall. A tangential force of 0.30N is applied parallel to the loaded face along the long direction of the block. The shear modulus of the rubber is 5.0×10^{5}Pa. Calculate the lateral displacement of the loaded face relative to the fixed face.

Chapter 2

TEMPERATURE AND HEAT

INTRODUCTION

It is a cold Miono morning and the school gate is closed. Kipanga arrives first, wrapped in a jumper two sizes too big, and touches the metal bars to push them open.

Kipanga: *"Aa! This gate is freezing! Somebody has left it in the fridge overnight."*

He stumbles backward, sits down on the wooden bench by the gate, and lets out a theatrical sigh.

Kipanga: *"But this bench, this bench is warm. I should have come and sat here first. The gate is definitely colder than the bench. Anyone who disagrees is a fool."*

Kipute, who has been standing three steps away with her hands in her pockets, produces a small thermometer from her bag and holds it first against the metal gate, then against the wooden bench. She looks at it for ten seconds in silence, then holds it up for Kipanga to read.

Kipute: *"Fourteen degrees. Both of them. Same reading to the nearest tenth."*

Kipanga stares at the thermometer. He stares at the gate. He stares at the bench. He stares at Kipute. He stares at the thermometer again.

Kipanga: *"The thermometer is broken."*

Kipute: *"The thermometer is not broken."*

Kipanga: *"Then my hand is broken."*

Kipute: *"Your hand is fine, Kipanga. It is your understanding of what 'cold' means that is broken."*

At that moment Mr. Akilikubwa arrives, carrying his briefcase in one hand and his morning tea in the other. He has been listening from the corner since the word "fridge". He sets down his briefcase, looks at the gate, looks at the bench, looks at Kipanga sitting on the bench and still squinting at the thermometer, and smiles the smile of a teacher who has just been handed a two-week chapter on a silver platter.

Mr. Akilikubwa: *"Congratulations, Kipanga. You have just discovered why thermometers were invented, why your fingers are not a reliable substitute for them, and why the next fifty pages of your physics book exist. Your hand is not broken and the thermometer is not broken. The thing flowing between your hand and the metal gate is not the same thing the thermometer is reading, and sorting those two things out is the work of this entire chapter. Come inside. The bench can wait."*

Kipanga follows him in, still holding the thermometer, still looking at the gate as if it has personally insulted him. Somewhere in the back of his mind, a small voice is asking whether the bench has been lying to him all these years too.

Chapter 1 taught us that solids change their length when they get hot. The SGR railway between Dodoma and Dar is still expanding as you read this sentence, a thermostat is clicking somewhere in a fridge, and a brass ring somewhere is grudgingly fitting over a steel shaft it was heated to accept. In all of these, something called **temperature** is driving the show, and yet we have been using the word without ever asking what it really means. *What exactly is temperature? How is it different from heat? How do we measure it reliably when our own fingers, as Kipanga has just discovered, cannot be trusted?* And once we can measure it, what can we calculate, how much heat does it take to boil a kettle, to melt a block of ice, to cool a hospital ward, to freeze a lake from the top down?

This chapter answers those questions in order. The opening sections build the machinery of temperature. They explain what temperature is, how the Zeroth Law gives thermometers their logical foundation, what counts as a thermometric property, how different thermometers are constructed, and why the thermodynamic scale is the only one that does not depend on the liquid inside the glass. The later sections then switch on the calculator and develop specific heat capacity, latent heat, heating curves, the ice on a pond, and the refrigerator in the kitchen. By the end, you will know why metal feels colder than wood at the same temperature, although Kipanga will have to wait for Chapter 3 for the full answer, why a kettle takes longer to boil than an equivalent mass of ice takes to melt, and why your own body's sweating is one of the most efficient cooling systems on the planet.

Welcome to Chapter 2. Your fingers have been lying to you all your life, and it is time to meet the instruments that tell the truth.

THE CONCEPT OF TEMPERATURE AND THE ZEROTH LAW

Temperature is not heat

Before we can measure temperature, we have to know what it is, and what it is not. The two most common mistakes students make with thermal physics are both versions of the same confusion, treating temperature and heat as if they were the same thing. They are not, and the distinction is worth getting right before another paragraph passes.

Temperature *is a scalar measure of the degree of hotness of a body*. It tells us *how hot* the body is, but it says nothing by itself about *how much energy* the body contains. A teaspoon of boiling water and a swimming pool of boiling water are at exactly the same temperature (100 °C at standard pressure), but the energy contained in the pool is millions of times greater than the energy contained in the teaspoon. If you spilled the teaspoon on your hand, you would yelp. If you fell into the swimming pool, the consequences would be considerably less comic. Same temperature, vastly different amounts of stored energy.

Heat, by contrast, *is a flow of energy from one body to another because of a temperature difference*. Heat is never at rest. It is always going somewhere, from the hotter body to the cooler body, until the two bodies are at the same temperature and the flow stops. The gate in the morning story was not transferring its "coldness" to Kipanga's hand; it was accepting heat **from** Kipanga's hand (because Kipanga at 37 °C was hotter than the gate at 14 °C), and **Kipanga's hand felt the loss as a sensation of cold**. The wooden bench was accepting heat from Kipanga's trousers at a much slower rate, which is why the bench felt warmer. Both surfaces were at the same temperature. The thermometer, being a more honest judge of hotness than Kipanga's skin, reported the truth.

Thermal equilibrium

Two bodies placed in contact with each other exchange heat until they reach a state in which there is no further net flow of energy between them. At that point, we say the two bodies are in **thermal equilibrium**. The thermometer pressed against the gate reached thermal equilibrium with the gate after a few seconds, then the same thermometer pressed against the bench reached thermal equilibrium with the bench a few seconds later. Both times it showed 14 °C, because both times it arrived at the same final state.

Thermal equilibrium has a simple physical meaning. The hotter body loses energy until its hotness drops, the colder body gains energy until its hotness rises, and the two meet in the middle. The process ends when neither body has any more reason to give or receive. Whatever quantity it is that determines this "meeting in the middle" is what we call temperature, and two bodies at the same temperature are defined as being in thermal equilibrium with each other.

The Zeroth Law of Thermodynamics

Now here is the beautiful statement that makes thermometers logically possible. Suppose Body A is in thermal equilibrium with Body C, and Body B is, separately, also in thermal equilibrium with Body C. A and B have never been in contact with each other, not even briefly. Can we conclude anything about A and B?

The **Zeroth Law of Thermodynamics** answers: yes, we can. *If A is in thermal equilibrium with C, and B is in thermal equilibrium with C, then A and B are also in thermal equilibrium with each other, and therefore at the same temperature.*

This sounds almost obvious. It is not. It is an experimental law of nature that happens to be true for the universe we live in, but there is no logical reason it had to be this way. The Zeroth Law is what guarantees that temperature is a well-defined, transitive property of bodies, rather than a relationship that only makes sense two bodies at a time. It is the law that justifies the existence of thermometers.

To see why, think about what a thermometer actually does. A thermometer is a small body (the "C" of our argument) that we bring into contact with the body we want to measure (the "A") until the thermometer and A reach thermal equilibrium. We read the thermometer's scale and record a number, say 14 °C. Later, we bring the same thermometer (having cooled down) into contact with another body B, reach thermal equilibrium, and again read 14 °C. The Zeroth Law now lets us conclude, without ever bringing A and B into contact, that A and B are in thermal equilibrium with each other, that is, at the same temperature. Without the Zeroth Law, this reasoning would not be valid, and a thermometer would not be a meaningful instrument.

Why is it called the "zeroth" law?

Historically, the First and Second Laws of Thermodynamics were formulated and named before the Zeroth Law. When physicists later realised that the transitivity of thermal equilibrium was more fundamental than either of them, because the First and Second Laws both assume you already know what temperature is, they faced an awkward naming problem. Renumbering the First and Second Laws would have broken a century of textbook references. Ralph Fowler proposed the solution in the 1930s: give the new law the number zero, placing it logically before the First Law without renumbering anything. The name stuck. The Zeroth Law sits at the foundation of thermodynamics and is called "zeroth" for the same reason a physicist starts counting rows at zero rather than one, because sometimes the thing that comes first is best labelled as coming before the thing we used to call first.

Theory and you have been heating each other up for a while. Time for a little thermal equilibrium, in the form of two worked examples.

BINDER Example 1

Two metal blocks X and Y are brought into contact with each other and left until they reach thermal equilibrium. The two blocks are then separated, and block X is brought into contact with a third, cooler block Z, and the two are left until they reach

thermal equilibrium. (a) What can you conclude about the relative temperatures of Y and Z without placing Y and Z in contact? (b) Which physical principle allows you to make this prediction? (c) Explain how the same principle guarantees that a thermometer is a logically valid instrument for comparing the temperatures of two different bodies.

Solution

(a) Block Y is at a higher temperature than block Z. Blocks X and Y reached thermal equilibrium first, so they share a common temperature, call it T_{XY}. Block X was then brought into contact with the cooler block Z and reached a new thermal equilibrium at some temperature T_{XZ} lower than T_{XY}, because X was the hotter of the two and lost heat to Z. Therefore Y, still at T_{XY}, is hotter than Z, which is now at $T_{XZ} < T_{XY}$. We reach this conclusion without Y and Z ever having touched each other.

(b) The principle is the **Zeroth Law of Thermodynamics**, which states that *if two bodies are each in thermal equilibrium with a third body, they are in thermal equilibrium with each other.* It is the Zeroth Law that allows us to use X as a common reference in the comparison between Y and Z.

(c) A thermometer is a body that is brought into thermal equilibrium first with one object and then with another. The Zeroth Law guarantees that two objects giving the same thermometer reading must be at the same temperature, and that two objects giving different readings are at different temperatures. Without this law, the thermometer's reading for object A would tell us nothing about how object A compares with object B, and the instrument would have no measuring power. The Zeroth Law is therefore not just a footnote in the theory of heat, it is the law that makes thermometry possible at all.

Making Sense of the Answer: *The argument works because temperature is transitive: if X shares a temperature with Y, and later with Z, then the three temperatures can be compared by using X as a go-between. Notice that we never needed to quote a number for any of the three temperatures. The physical ordering was enough. That is also how any real thermometer works when we compare two rooms, two patients, or two buckets of water without ever bringing them into the same place.*

Think Like a Physicist: *Whenever you are asked about temperatures of bodies that have been in contact in different sequences, draw a chain: A equilibrates with B, B equilibrates with C, and so on. The Zeroth Law lets you follow the chain with confidence that every link preserves equality of temperature. The only way the chain breaks is if a body is heated or cooled between contacts, which will always be stated in the question.*

REAL Example 2

Mr. Akilikubwa is on break duty in the staff room when Kipanga appears at the door, looking frustrated.

Kipanga: *"Sir, I have just been to the clinic because I thought I had a fever, and the nurse put a thermometer in my mouth. I thought she was going to leave it there forever. Why did she not read it immediately? I know my own body, and if I am feeling hot, I am hot. A thermometer is just a piece of glass. It measures its own temperature, not mine. The moment she put it in my mouth, it started cooling down my tongue, so it was obviously going to read lower than I really was."*

Mr. Akilikubwa invites Kipanga to sit down and explain to Kipanga: (a) why the nurse waited before taking the reading, (b) why the thermometer's reading eventually gave the correct body temperature despite being, as Kipanga correctly noted, "just a piece of glass", and (c) whether Kipanga's worry that the thermometer cooled his tongue is actually a problem for the measurement.

What was Mr. Akilikubwa's response?

Solution

(a) The nurse waited because the thermometer and Kipanga's mouth were not initially in thermal equilibrium. When the thermometer was first placed in Kipanga's mouth, it was at room temperature (about 25 °C) and his mouth was at body temperature (about 37 °C). Heat began to flow from the mouth to the thermometer, and the thermometer's mercury column rose as its temperature climbed. Until thermal equilibrium had been reached, the thermometer was reading its own current temperature, which was somewhere between 25 °C and 37 °C. Reading it immediately would have given a value below the true body temperature.

(b) Kipanga is correct that the thermometer measures its own temperature. But the Zeroth Law guarantees that once the thermometer has reached thermal equilibrium with his mouth, the thermometer's temperature equals the mouth's temperature. At that moment, "the thermometer's temperature" and "Kipanga's body temperature" are numerically the same value, and reading the thermometer's scale reveals the body temperature.

(c) Kipanga's concern is real but small. The thermometer does indeed cool his tongue slightly, by drawing a little heat from his mouth on its way to equilibrium. But the amount of heat required to warm the thermometer's small mercury column from 25 °C to 37 °C is tiny compared with the heat content of his whole mouth and the blood flowing through it, so the mouth's temperature drops by a negligible fraction of a degree. The thermometer essentially reaches the mouth's original temperature, not the other way round. This is one of the reasons clinical thermometers are made with very small mercury bulbs: the smaller the thermometer's thermal capacity, the less it disturbs the body it is measuring.

We will meet thermal capacity formally later in this chapter; for now, it is enough to know that the nurse's instrument is designed to read the patient, not to change the patient.

Making Sense of the Answer: *Kipanga's three objections each contained a grain of truth, and Mr. Akilikubwa's answers did not dismiss any of them. The thermometer really does measure its own temperature, it really does cool the tongue slightly, and the nurse really did have to wait. What Kipanga missed was that thermal equilibrium resolves all three worries at once. Once the thermometer and the mouth are at the same temperature, there is nothing left to cool, nothing left to wait for, and the two temperatures are numerically equal. The Zeroth Law quietly guarantees that the single number on the thermometer is also the number for Kipanga's body.*

Think Like a Physicist: *Whenever a measurement involves putting a small instrument in contact with a larger system, ask two questions. First, how long until thermal equilibrium? That tells you how long to wait before reading. Second, how much does the instrument disturb the system it is measuring? That tells you how large your instrument is allowed to be. Clinical thermometers have tiny bulbs for exactly this reason. Laboratory calorimeters have water jackets for the same reason. Every good measuring instrument is a compromise between reading quickly and reading honestly.*

With the concept of temperature secured and the Zeroth Law standing behind every thermometer ever made, we are ready to ask a more practical question. *Any physical property that changes with temperature can in principle be used to build a thermometer, but what makes one property good and another useless?* The next section separates the contenders from the pretenders.

THERMOMETRIC PROPERTIES AND EMPIRICAL TEMPERATURE SCALES

The Zeroth Law tells us that thermometers are logically possible. It does not yet tell us how to build one. A working thermometer needs some physical quantity inside it that changes in a trustworthy way as temperature changes, and it needs a scale printed alongside that quantity so that a reader can turn the change into a number. This section is about what a trustworthy physical quantity looks like, and about how a simple scale is built from it by choosing two reference points and drawing a straight line between them.

Thermometric properties

A **thermometric property** *is any physical property of a substance that varies with temperature in a way that we can read off and turn into a temperature measurement.* In principle, almost any physical quantity that depends on temperature could be used. This include:

- the length of a metal rod,
- the volume of a liquid,
- the pressure of a gas held at constant volume,
- the electrical resistance of a wire,
- the electromotive force generated at the junction of two dissimilar metals,
- the colour of a heated piece of iron.

The whole of thermometry is an exercise in choosing among these candidates the ones that are reliable and practical, and discarding the ones that look attractive until you try to use them.

Three characteristics separate a good thermometric property from a bad one.

1) The property must vary continuously with temperature

A property that changes smoothly as the temperature rises gives a single, unambiguous reading at every temperature. A property that jumps suddenly at a particular point, or that has flat regions where it refuses to change, is useless as a thermometer because it cannot distinguish between the temperatures in those flat regions. The length of a mercury column meets this test; the colour of a metal rod (which is black at room temperature and stays black until several hundred degrees) **fails** it in the lower range.

2) The property must be measurable with sufficient precision

A thermometer is only as accurate as the reading we can take off its scale. If a property changes by a huge amount for a given temperature change, the thermometer is sensitive and a small temperature difference can be resolved. If the property changes by only a tiny amount, the thermometer is insensitive and small temperature differences go undetected. The pressure of a gas held at constant volume changes substantially with temperature and is easy to read off a manometer. The speed of sound in a solid also changes with temperature, but only by a fraction of a percent per degree; using it as a thermometric property would demand extraordinary equipment to detect each degree.

3) The property must be reproducible

The same substance, brought to the same temperature under the same conditions, must give the same value of the property every time. A substance that behaves differently on Monday from how it behaves on Tuesday, or that gives different readings in different laboratories, cannot anchor a scientific instrument. Pure mercury in a sealed capillary is fully reproducible. The colour of a roast chicken at a given temperature is not, for a great many reasons.

A property that passes all three tests can be used to build a thermometer. A property that fails **any** one of them cannot. Most of the candidates listed earlier (mercury volume, gas pressure, platinum resistance, thermocouple electromotive force) pass all three and are exploited in real instruments. Some others (length of a metal rod, specific heat, thermal conductivity) pass in principle but are too insensitive or too inconvenient to build a practical thermometer around. And a few (colour of a heated rod, rate of sound travel) fail outright in one way or another.

Building a temperature scale

Even a trustworthy thermometric property is useless on its own. The property gives us a number, such as "the mercury column is 43 mm long", or "the gas pressure is 1.27 times its ice-point value", and by itself this number is not a temperature. To convert the property into a temperature we need a **scale:** *a rule that assigns a temperature value to every possible reading of the property*.

The simplest rule is a linear one. Pick two fixed points at which the temperature is known, measure the thermometric property at each of these points, and assume the property varies linearly with temperature between them. Any reading in between can then be converted into a temperature by the straight-line formula.

The Celsius scale is built in exactly this way. Its two fixed points are the **ice point**, defined as *the temperature of pure melting ice at standard atmospheric pressure and assigned the value* $0°C$, and the **steam point**, defined as *the temperature of pure water boiling at standard atmospheric pressure and assigned the value* $100°C$.

Let X_0 be the value of the thermometric property at the ice point, X_{100} its value at the steam point, and X_θ its value at an unknown Celsius temperature θ.

The assumption of linearity between the two fixed points means that a graph of the property X against the temperature θ is a straight line passing through the two calibration points $(0, X_0)$ and $(100, X_{100})$. On a straight line, any two segments between the endpoints have the same slope. Taking the segment from the ice point up to the unknown temperature and comparing it with the segment from the ice point all the way up to the steam point:

$$\frac{X_\theta - X_0}{\theta - 0} = \frac{X_{100} - X_0}{100 - 0}$$

Both sides of this equation are just the slope of the line X versus θ, computed two different ways, and linearity guarantees they are equal. Simplifying the denominators:

$$\frac{X_\theta - X_0}{\theta} = \frac{X_{100} - X_0}{100}$$

Cross-multiplying to clear the fractions:

$$100\,(X_\theta - X_0) = \theta\,(X_{100} - X_0)$$

Dividing both sides by $(X_{100} - X_0)$ to isolate θ:

$$\theta = \frac{100\,(X_\theta - X_0)}{X_{100} - X_0}$$

Writing this in the more compact form:

$$\boldsymbol{\theta = \left(\frac{X_\theta - X_0}{X_{100} - X_0}\right) \times 100°C}$$

This is the **interpolation formula** for the Celsius scale, and it lies behind every empirical thermometer calibrated on ice and steam. Substituting $X = \ell$ (length of the mercury column) gives the Celsius temperature from a mercury-in-glass thermometer; substituting $X = P$ (gas pressure at constant volume) gives it from a gas thermometer; substituting $X = R$ (electrical resistance) gives it from a platinum resistance thermometer. The formula does not care what the property is. It only assumes that the property varies linearly with temperature between the two fixed points.

Why different thermometers disagree between fixed points

Here is the catch. The interpolation formula **assumes** linearity. It does not guarantee it. Real thermometric properties vary **approximately** linearly over short temperature ranges, but not exactly, and the departure from linearity is different for different substances. The mercury column does not expand with perfect proportionality to temperature; neither does the gas pressure, neither does the platinum resistance. Each of them follows its own slightly curved graph of property against temperature, and each of them is forced by calibration to agree at the two fixed points.

Between those two fixed points, however, the three curves pass through different regions, and the three thermometers give three different readings for the same physical temperature. A mercury thermometer and a platinum resistance thermometer, both calibrated at $0°C$ and $100°C$, might both read $49.3°C$ and $50.2°C$ for a water bath that a gas thermometer reads as $50.0°C$. None of the three is broken. Each is simply reporting what its own thermometric property tells it, on the assumption that the property varies linearly, when in fact the property's graph has a slight curvature that the formula cannot see. The thermometers agree at $0°C$ and $100°C$ by design. Between them, they disagree because the physics of their properties disagrees.

This is not a failure of the interpolation method. It is a feature of empirical temperature scales. An **empirical scale** *is one that depends on the behaviour of a particular thermometric substance, and different substances give different scales.* Later in this

chapter we will meet a way around this problem: the thermodynamic scale, which is independent of the substance used in the thermometer and to which all real thermometers can be compared. For the moment, the lesson is this. Two thermometers calibrated at the same fixed points will agree at those fixed points exactly and at every other temperature only approximately, and *the size of the disagreement reveals the* ***non-linearity*** *of one or both of their thermometric properties.*

Faulty thermometers

There is one more quick application of the interpolation formula before we leave this section. Suppose a mercury-in-glass thermometer has been manufactured carelessly, and its "0°C" mark has been engraved where the mercury actually sits at 4°C, and its "100°C" mark has been engraved where the mercury actually sits at 98°C. The mercury column itself is blameless. It still expands by the same amount per degree as it always did, because its physics has not changed. What has changed is only the labels painted on the outside of the glass, and the labels are wrong by a fixed amount at each end. *If the faulty thermometer reads θ_f, what is the true Celsius temperature?*

The key insight is that the faulty reading θ_f is itself a valid thermometric property of the liquid being measured. It varies continuously with temperature, it can be read off the scale with precision, and it is reproducible from one measurement to the next (it is, after all, simply a misprinted mercury column). So we can treat the faulty reading exactly as we would treat any other thermometric property X, and apply the interpolation formula with the faulty readings as the X-values.

To use the formula, we need three numbers: the faulty thermometer's reading at the true ice point, at the true steam point, and at the unknown temperature. Let θ_0^f be the faulty reading when the thermometer is immersed in pure melting ice (at a true temperature of 0°C), and θ_{100}^f be the faulty reading when the thermometer is immersed in boiling water (at a true temperature of 100°C). By construction, θ_0^f and θ_{100}^f are simply the true temperatures at which the faulty "0°C" and "100°C" marks have been engraved. In our example above, $\theta_0^f = 4°C$ and $\theta_{100}^f = 98°C$.

Substituting $X_\theta \to \theta_f$, $X_0 \to \theta_0^f$, and $X_{100} \to \theta_{100}^f$ into the interpolation formula gives the true temperature θ from the faulty reading:

$$\boldsymbol{\theta = \left(\frac{\theta_f - \theta_0^f}{\theta_{100}^f - \theta_0^f}\right) \times 100^\circ C}$$

The formula is identical in structure to the calibration formula; only the labels on the quantities have changed. A faulty thermometer is just an ordinary interpolation problem in disguise.

Theory and you have been heating each other up again. Let us reach thermal equilibrium with three worked examples before moving on to the zoo of real thermometers in the next section.

BINDER Example 3

A student proposes using the colour intensity of a glowing metal rod as a thermometric property, arguing that the rod goes from dull red to bright red to orange to white as the temperature rises. Evaluate this proposal by checking it against the three characteristics of a good thermometric property. Compare your conclusions with the corresponding properties of the electrical resistance of a platinum wire.

Solution

To evaluate this, we must answer three crucial questions:

1) Does the colour intensity vary continuously with temperature across the useful range?

The colour intensity of a heated rod varies continuously with temperature only above a threshold of roughly 500°C, below which the rod emits no visible light at all and appears simply as dark metal. Below this threshold, the property is not just insensitive, it is absent: the rod is black at 100°C, black at 200°C, and black at 400°C. A thermometer based on colour intensity would be blind across this entire range. The electrical resistance of a platinum wire, by contrast, varies smoothly from below room temperature upward with no threshold, and a single wire can be used from below the ice point to several hundred degrees above the steam point.

2) Can it be measured with sufficient precision at every temperature in the range?

Colour intensity is hard to measure precisely because human eyes disagree with each other and because the perceived colour depends on ambient lighting, on the angle from which the rod is viewed, and on the chemical composition of the rod's surface. Instruments called optical pyrometers do exist for measuring colour-temperature in industrial settings, but their precision even at best is several degrees, and at low temperatures they cannot operate at all. A platinum resistance measurement, by contrast, can resolve a hundredth of a degree with a Wheatstone bridge and is limited only by the precision of the electrical instruments, not by the human observer.

3) Is it reproducible across different samples and different laboratories?

Reproducibility across samples is poor for colour intensity because different metals emit different spectra at the same temperature, and the same metal with a tarnished or oxidised surface emits differently from a freshly polished sample. The same piece of iron heated to 800°C in Dar es Salaam and in Dodoma would be reported with slightly different colour readings

by identical optical pyrometers if the surfaces had weathered differently. Platinum resistance is fully reproducible because high-purity platinum has a resistance-temperature relationship that is stable across samples to better than a part in ten thousand and is not affected by surface tarnishing.

On all three counts, colour intensity fails and platinum resistance succeeds. Colour intensity can be used as a rough qualitative indicator in a blacksmith's forge, but not as a scientific thermometer. Platinum resistance is the standard by which all other precision thermometers are compared.

Making Sense of the Answer: *The three characteristics we checked are not a bureaucratic list. Each of them rules out a different kind of failure. Continuity rules out thermometers that can only read in part of the range, precision rules out thermometers that cannot resolve small differences even where they do read, and reproducibility rules out thermometers that give different answers in different laboratories. A property can pass two tests and fail the third, in which case it is still unsuitable. Colour intensity fails all three at once, which is why blacksmiths use it for judgement and physicists use platinum.*

Think Like a Physicist: *When somebody proposes a new thermometric property, ask the three questions in order. Does the property change across the range you care about? Can you measure the change precisely? Does the same substance give the same reading every time? If the answer to any one of the three is no, walk away, however clever the idea might otherwise be. The graveyard of physics is full of beautiful thermometers that fail one of these tests in an inconvenient place.*

BINDER Example 4

A mercury-in-glass thermometer has been manufactured carelessly. Its "0°C" mark has been engraved where the mercury actually sits at 3°C, and its "100°C" mark has been engraved where the mercury actually sits at 97°C. A student uses this thermometer to measure the temperature of a warm liquid, and the thermometer reads 59°C. What is the true Celsius temperature of the liquid?

Solution

The mercury column itself expands linearly between the two physical endpoints, which are at true temperatures of 3°C and 97°C. The engravings on the glass divide this physical interval into 100 equal parts, but the parts span only $97 - 3 = 94$ true degrees rather than the nominal 100 degrees. The thermometer therefore under-reads at low temperatures and over-reads at high temperatures, and a single reading in the middle of the scale must be corrected.
Applying the interpolation formula with the faulty readings as the X-values:

$$\theta = \left(\frac{\theta_f - \theta_0^f}{\theta_{100}^f - \theta_0^f}\right) \times 100°C$$

With $\theta_f = 59°C$, $\theta_0^f = 3°C$, and $\theta_{100}^f = 97°C$:

$$\theta = \left(\frac{59 - 3}{97 - 3}\right) \times 100°C = 59.57°C$$

The true temperature of the liquid is approximately 59.6°C, slightly higher than the faulty thermometer suggested.

Making Sense of the Answer: *The correction of* 0.6°C *looks small, and in a kitchen it would not matter. In a chemistry laboratory where a reaction is controlled to within one degree, a faulty thermometer that under-reads by half a degree at* 60°C *is an expensive piece of glass. The lesson is that calibration errors are not romantic; they accumulate quietly into bad science. A good physics student, meeting a new thermometer, always measures the ice and steam points before trusting anything else it says.*

Think Like a Physicist: *Whenever a thermometer is suspected of being faulty, ask yourself what its readings at the true ice and true steam points really are. Those two readings, substituted as* θ_0^f *and* θ_{100}^f *in the interpolation formula, convert any other reading into its true value. The same formula works whether the mistake is in the engraving, in the mercury mass, or in the glass capillary bore. Empirical interpolation is remarkably forgiving of manufacturing errors, provided you know what the errors are.*

REAL Example 5

A mercury-in-glass thermometer and a platinum resistance thermometer are both carefully calibrated at the ice point (0°C) and the steam point (100°C). A student immerses both thermometers in a water bath that a precision gas thermometer reports to be at exactly 50.0°C, waits for thermal equilibrium, and reads the two thermometers. The mercury thermometer reads 49.3°C and the platinum resistance thermometer reads 50.2°C. The student exclaims that one of the thermometers must be broken, since they were calibrated identically. Evaluate this claim and explain what the disagreement actually reveals about the two thermometric properties.

Solution

Neither thermometer is broken. The student has overlooked the single assumption on which the interpolation formula rests: the thermometric property varies **linearly** with temperature. In reality, the variation is only approximately linear, and the departure from linearity is different for different properties. The three readings in the question are enough to reveal what is happening.

Take the gas thermometer's reading of 50.0°C as the true temperature, since an ideal gas at constant volume is known to vary very nearly linearly with temperature. The mercury thermometer, applying the interpolation formula to its own measured column length, returns 49.3°C, which is 0.7°C below the true value. The formula does its arithmetic faithfully, so the under-reading can only be produced by the mercury column at 50°C being slightly shorter than a linear expansion between the two fixed points would predict. This means the mercury's volume expansion falls **below** a straight-line relationship in the middle of the interval. The platinum thermometer, by the same argument, returns 50.2°C, which is 0.2°C above the true value. Its resistance at 50°C is therefore slightly greater than a linear rise between the fixed points would predict, which means platinum resistance rises **above** a straight-line relationship in the middle of the interval.

The disagreement between the two thermometers therefore reveals that the two thermometric properties are **non-linear in opposite directions:** mercury volume expansion is slightly **sub-linear** between the fixed points, while platinum electrical resistance is slightly **super-linear.** The mercury's non-linearity is also larger in magnitude (0.7°C versus 0.2°C), which is consistent with the fact that platinum resistance is widely used as a sub-standard thermometer while mercury is used only for routine work.

Making Sense of the Answer: *Two thermometers calibrated at the same fixed points must agree at those fixed points by design, but they are under no obligation to agree at any other temperature. The student's reaction is a common one: if two instruments disagree, one of them must be lying. A more experienced physicist would reach a different conclusion: if two instruments disagree and neither is broken, the disagreement is telling you something about the instruments themselves, not about the thing they are both measuring.*

Think Like a Physicist: *Whenever you calibrate a thermometer on two fixed points, you are promising the reader only that the thermometer will agree with itself at those two points. Between them, the reading depends on the shape of the property-versus-temperature curve, which is different for every thermometric property. Two thermometers based on different properties will disagree slightly between fixed points in a way that no amount of careful calibration can fix. The only way to eliminate the disagreement is to escape empirical scales altogether, which is what the thermodynamic scale was invented to do.*

We now have the two ingredients that every working thermometer needs: a reliable physical property to respond to temperature, and a simple scale to turn the property's reading into a number. But we have been speaking of thermometers in the abstract, as though any thermometer will do. In real laboratories and real hospitals, different thermometers are chosen for different jobs, and the differences between them are not cosmetic. The mercury thermometer, the gas thermometer, the platinum resistance thermometer, and the thermocouple have very different strengths and very different weaknesses, and knowing which is which can be the difference between a measurement that matters and a number that misleads. It is time to meet them individually.

TYPES OF THERMOMETERS

The previous section gave us the formula that every empirical thermometer uses to turn a physical property into a temperature. What it did not give us is the thermometers themselves. Different jobs demand different thermometers, and the choice between a mercury bulb, a platinum coil, a thermocouple, and an optical pyrometer is not a matter of taste. Each instrument has a range in which it works well, a range in which it works poorly, and a range in which it does not work at all. The purpose of this section is to meet each of them in turn, see how it is built, see what it is good for, and see where it fails.

1. Liquid-in-glass thermometers

The liquid-in-glass thermometer is the one most students have actually held in their hand. It consists of a thin-walled glass bulb at the bottom, a fine capillary tube rising from the bulb, and a graduated scale etched or printed alongside the capillary. A small quantity of liquid fills the bulb and rises partway up the capillary. When the bulb is placed in contact with a warmer body, the liquid expands and the column rises up the capillary. When placed in a cooler body, the liquid contracts and the column falls. Th**e *length of the column* is the thermometric property.**

Mercury is the most common working liquid for laboratory and clinical use. Pure mercury expands uniformly over a very wide range of temperatures, from its freezing point at −39°C to its boiling point at 357°C, and across almost all of this range *it expands so nearly linearly* with temperature that the small non-linearities we met in the previous section contribute less than a degree of error over the whole span from ice point to steam point. Mercury is *opaque and silvery*, so the top of the column is *easy to see against the graduated scale*. It *does not wet the glass,* which means the *column rises and falls cleanly without leaving droplets behind* that would distort the reading.

Mercury's strengths come with three real limitations. **First,** mercury freezes at −39°C, so a mercury thermometer cannot be used at all in polar or cryogenic conditions. **Second,** mercury is extremely toxic if spilled or inhaled as vapour, which makes a broken mercury thermometer in a classroom a small disaster that requires specialist cleanup. **Third,** mercury is expensive, and a laboratory-grade mercury thermometer is one of the more costly pieces of glassware in a school stockroom.

Alcohol (usually ethanol coloured with a red dye to make the meniscus visible) is used where mercury cannot go. Alcohol freezes at −115°C, so *it remains fluid deep below the mercury cut-off*, and it is therefore the liquid of choice for thermometers used to measure atmospheric temperatures in cold regions or the temperatures inside domestic freezers. Its weaknesses are

the mirror image of mercury's strengths. Alcohol *expands less uniformly with temperature than mercury*, so alcohol thermometers are slightly less accurate across any given range. Alcohol also boils at 78°C, which puts a *strict upper limit* on the temperatures it can measure. And because ethanol wets glass, the column tends to leave a film behind when it falls, which introduces a small *readout error* that must be corrected for.

The clinical thermometer is a specialised mercury-in-glass instrument designed for measuring human body temperature. It has two features that distinguish it from a laboratory thermometer. **First,** its scale is restricted to a narrow window around body temperature, typically 35°C to 42°C, which allows each degree to be spread across several millimetres of capillary for easy reading. **Second,** and more importantly, the capillary has a **constriction** just above the bulb, a short section narrower than the rest of the tube. When the bulb is warmed during use, the mercury expands and is forced through this constriction into the main capillary. When the thermometer is withdrawn and begins to cool, the mercury in the bulb contracts and tries to pull the column back down, but surface tension at the constriction breaks the column at that point. The mercury above the constriction cannot return to the bulb, and stays in place until the nurse has had time to read the scale. To reset the thermometer, the nurse gives it a sharp downward flick, which uses centrifugal force to drive the upper mercury back through the constriction into the bulb.

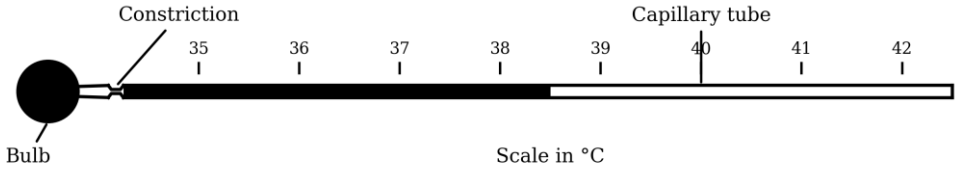


Figure: ***Clinical thermometer*** *showing the bulb, the constriction just above the bulb, and the narrow capillary tube carrying a scale calibrated for body temperature.*

2. Constant-volume gas thermometer

The constant-volume gas thermometer uses the ***pressure of a fixed mass of gas held at constant volume*** **as its thermometric property.** The gas is sealed inside a rigid glass bulb connected by a thin tube to a mercury manometer, and the temperature of the gas is changed by immersing the bulb in the body whose temperature is being measured. As the gas heats up, its pressure rises; the difference in mercury levels in the two arms of the manometer gives the pressure, and the pressure is substituted into the interpolation formula to give the temperature on the gas scale:

$$\theta_{\text{gas}} = \left(\frac{P_\theta - P_0}{P_{100} - P_0}\right) \times 100^\circ\text{C}$$

A sibling variant (**constant-pressure gas thermometer**) exists in which the pressure is held constant and the volume is allowed to change; the corresponding formula is $\theta_{\text{gas}} = \frac{V_\theta - V_0}{V_{100} - V_0} \times 100^\circ\text{C}$, but the constant-volume form is the one almost always used in practice.

The gas thermometer is often called the ***ideal*** **empirical thermometer**, because *the pressure of a dilute gas varies very nearly linearly with temperature over an enormous range*. The departure from linearity that we saw affecting the mercury and platinum thermometers is much smaller for a gas, and it becomes smaller still as the gas is made more and more dilute. In the limit where the amount of gas in the bulb tends to zero, the gas thermometer agrees exactly with the thermodynamic scale that we will meet in the next section. This is why the gas thermometer is used as the primary standard against which all other thermometers are calibrated.

The strength of the gas thermometer is its *high accuracy*. A carefully constructed laboratory gas thermometer can resolve temperatures to a thousandth of a degree. Its weaknesses are practical rather than physical. The apparatus is *large*, *fragile*, and *slow to respond*, because a large mass of gas takes time to reach thermal equilibrium with the body being measured. It is also unsuited to measuring rapidly changing temperatures, and it cannot be taken into the field. For these reasons the gas thermometer lives in the national metrology laboratory, not in the clinic or the classroom.

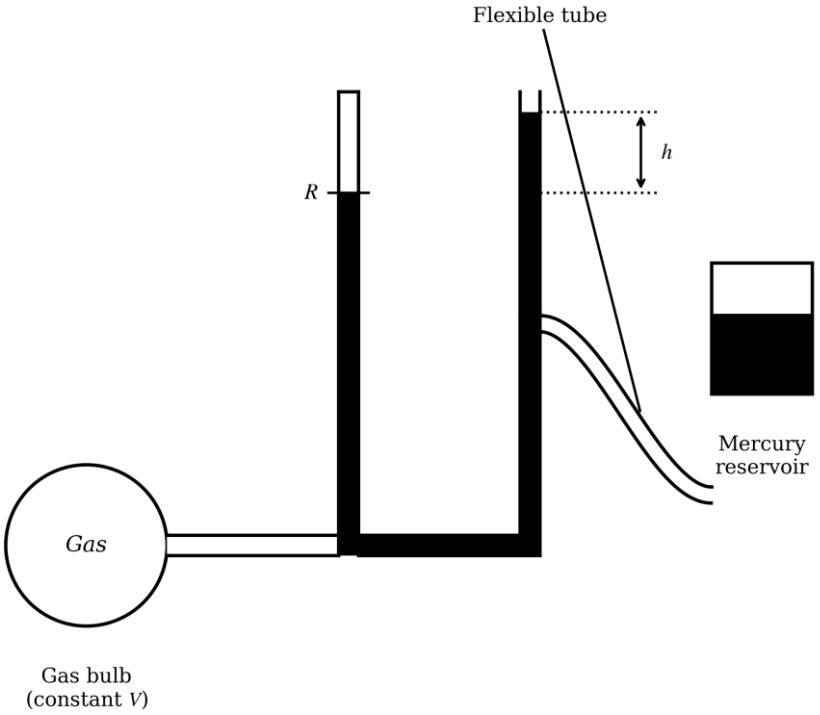


Figure: ***Constant-volume gas thermometer.*** *The gas is held in a rigid bulb whose temperature is to be measured. A flexible tube connects the right arm of the mercury manometer to a reservoir whose height is adjusted so that the mercury in the left arm always reaches the reference mark R, keeping the gas volume constant. The height difference h between the two arms is then a direct measure of the gas pressure.*

3. Platinum resistance thermometer

The platinum resistance thermometer uses the ***electrical resistance*** **of a pure platinum wire as its thermometric property.** Platinum is chosen because it has a *high melting point* (1773°C), is *chemically inert* across an enormous range of temperatures, *resists corrosion*, and has a *resistance that changes by a substantial and reproducible fraction* as the temperature changes. The wire is wound as a coil on a mica former to give a compact, mechanically stable element, and the coil is protected inside a quartz or ceramic sheath. The resistance is measured by connecting the coil as one arm of a Wheatstone bridge, with the other three arms provided by known resistors. At each temperature, the bridge is balanced and the coil's resistance is read off directly.

The resistance at an unknown temperature θ is substituted into the interpolation formula in the same way as the mercury column length or the gas pressure:

$$\theta_R = \left(\frac{R_\theta - R_0}{R_{100} - R_0}\right) \times 100°C$$

For calibration work that demands a functional form, the resistance of platinum as a function of true temperature is frequently written as $\mathbf{R_\theta = R_0(1 + a\theta + b\theta^2)}$, a parabolic form in which a and b are small constants determined by the particular wire. An alternative form is $\mathbf{R_\theta = R_0(1 + 800b\theta - b\theta^2)}$, which is algebraically equivalent to the first after a redefinition of a. The platinum resistance thermometer is also the thermometric instrument of choice for many sub-standard calibrations, and a variant called the **thermistor**, made from semiconducting material whose resistance **decreases** with temperature, is used for applications where small size and fast response matter more than absolute accuracy.

The platinum resistance thermometer's strengths are its *high accuracy* (second only to the gas thermometer), its *wide working range* (from below $-200°C$ to roughly $1000°C$), and its *stability* across years of use without recalibration. Its weaknesses are its *slow response time* (the coil has a larger thermal mass than a thermocouple junction), its *need for a stable current source and a bridge circuit*, and the fact that *self-heating* by the measuring current itself can introduce a small error at high precision.

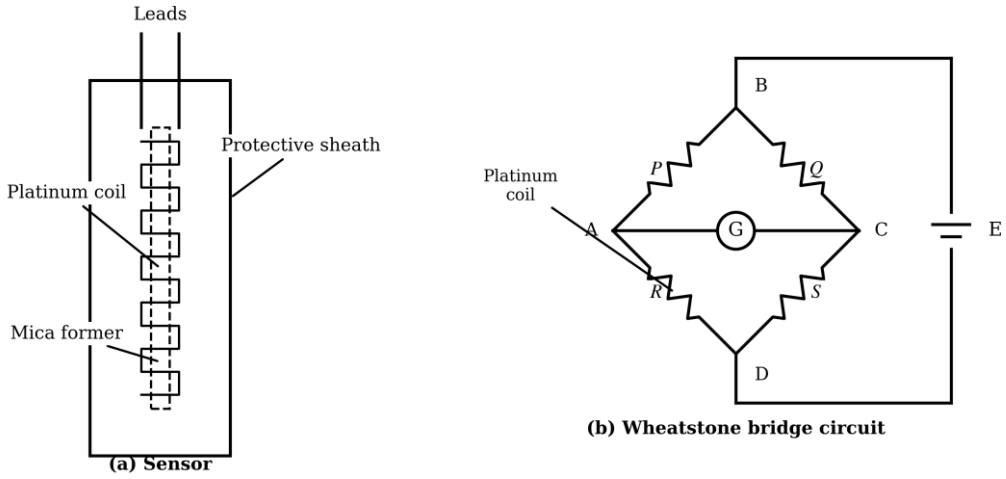


Figure: ***Platinum resistance thermometer.*** *The sensor shown in (a) consists of a platinum coil wound on a mica former and enclosed in a protective sheath. The measurement circuit shown in (b) connects the coil as one arm of a Wheatstone bridge, with known resistors P, Q, S forming the other three arms, a galvanometer G across the diagonal, and a battery E providing the current.*

4. Thermocouple thermometer

The thermocouple exploits a phenomenon discovered by Thomas Seebeck in 1821. When two wires of different metals are joined at both ends to form a closed loop and the two junctions are held at different temperatures, an electromotive force (EMF) appears around the loop, and a current flows if the loop is completed. This **Seebeck EMF** depends only on the pair of metals used and on the temperatures of the two junctions, not on the shape of the wires or the lengths between the junctions. A thermocouple thermometer consists of two such wires welded together at one end (the **hot junction**, placed in the body being measured) and connected to a high-resistance voltmeter at the other end (the voltmeter connections serve as the **cold junction**, held at a known reference temperature, usually 0°C in ice water). The voltmeter reads the ***Seebeck EMF,*** which is the **thermometric property.**

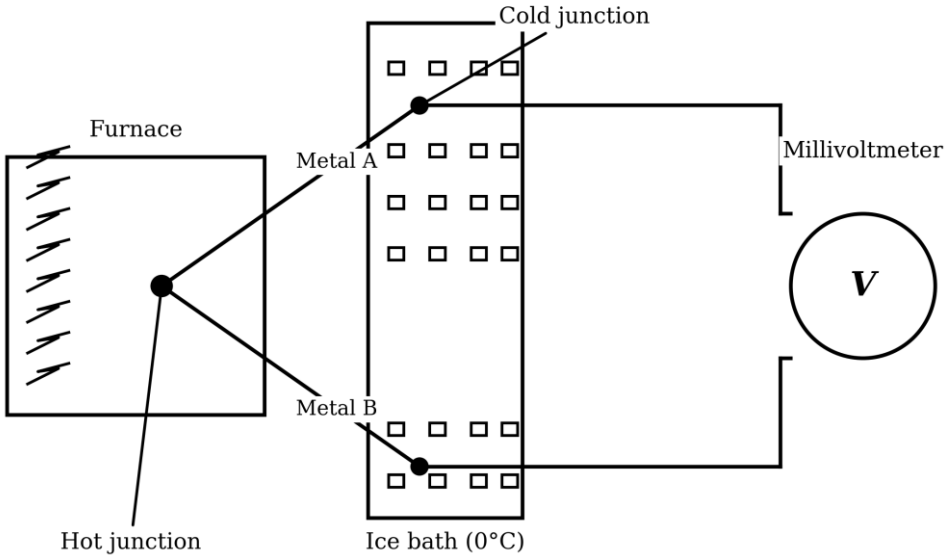


Figure: ***Thermocouple thermometer.*** *Two wires of dissimilar metals (Metal A and Metal B) are welded together at a hot junction placed inside the furnace whose temperature is being measured. A cold junction is formed where each wire meets a copper lead in an ice bath held at 0°C. The millivoltmeter reads the thermoelectric EMF generated around the loop.*

The Seebeck EMF is not a linear function of the hot-junction temperature. For a typical thermocouple with the cold junction at 0°C, the EMF varies with the hot-junction temperature θ according to a parabolic relation, which is commonly written in one of two algebraically equivalent forms:

$$\mathbf{E = a\theta + b\theta^2} \qquad \text{or} \qquad \mathbf{E = a\theta + \frac{1}{2}b\theta^2}$$

The two forms arise from different conventions in the literature, and a student should read any given problem carefully to see which form the constants a and b are fitted to. Both a and b are specific to the pair of metals used; a is always positive and b is always small and negative, which means the EMF rises rapidly with temperature at first, reaches a peak at some specific temperature, and then falls as the temperature continues to climb.

Like every other empirical thermometer, the thermocouple also has an interpolation formula built from the general form of the previous section by substituting $X = E$:

$$\theta_{tc} = \left(\frac{E_\theta - E_0}{E_{100} - E_0}\right) \times 100°C$$

When the cold junction is held at 0°C, E_0 is zero by construction, and the formula reduces to:

$$\theta_{tc} = \left(\frac{E_\theta}{E_{100}}\right) \times 100°C$$

This is the formula a student uses when given tabulated EMF readings at the fixed points and an EMF reading at the unknown temperature. The Seebeck parabolic equation introduced above, by contrast, describes the underlying non-linear physics of the thermocouple. The two formulas serve different purposes: the interpolation formula gives the thermocouple-scale reading a calibrated instrument would display, while the parabolic Seebeck equation gives the true EMF produced by the physics of the junction. A student must always check which formula a question is asking for before substituting numbers.

The temperature at which the EMF peaks is called the **neutral temperature**, and is denoted $\boldsymbol{\theta_N}$. Physically, the **neutral temperature** *is the hot-junction temperature at which the rate of increase of EMF with temperature has fallen to zero.* Beyond the neutral temperature, the EMF actually decreases as the hot junction gets hotter, so the same value of E can correspond to two different hot-junction temperatures: one below θ_N and one above. A thermocouple used at or near its neutral temperature is therefore ambiguous, and a thermocouple used beyond its neutral temperature is useless as a measuring instrument unless other information pins down which side of the peak the hot junction is on.

To find the neutral temperature from the first form, we differentiate E with respect to θ and set the result to zero, because a smooth curve is flat at its peak.

Differentiating $E = a\theta + b\theta^2$ with respect to θ:

$$\frac{dE}{d\theta} = a + 2b\theta$$

Setting $\frac{dE}{d\theta} = 0$ at the neutral temperature θ_N:

$$a + 2b\theta_N = 0$$

Solving for θ_N:

$$\boldsymbol{\theta_N = -\frac{a}{2b}}$$

The maximum EMF is the value of E at $\theta = \theta_N$. Substituting $\theta_N = -\frac{a}{2b}$ back into $E = a\theta + b\theta^2$:

$$\mathbf{E_{max}} = a\left(-\frac{a}{2b}\right) + b\left(-\frac{a}{2b}\right)^2 = -\frac{a^2}{2b} + \frac{a^2}{4b} = \boldsymbol{-\frac{a^2}{4b}}$$

The minus sign in $E_{max} = -\frac{a^2}{4b}$ does not mean the EMF itself is negative; it cancels against the negative value of b to give a positive maximum. Because $b < 0$, the quantity $-\frac{a^2}{4b}$ is positive.

For the second form of the Seebeck equation, $E = a\theta + \frac{1}{2}b\theta^2$, the same procedure gives:

$$\frac{dE}{d\theta} = a + b\theta, \quad \Rightarrow \quad \boldsymbol{\theta_N = -\frac{a}{b}}, \quad \Rightarrow \quad \mathbf{E_{max}} = \boldsymbol{-\frac{a^2}{2b}}$$

The result suggests that the neutral temperature is twice as large in the second form, and the maximum EMF is twice as large. However, this is not a physical difference between two different thermocouples; it is the same thermocouple described by two different constants fitted to two different forms of the same parabolic relationship. A student must always check which form a question uses before substituting numbers.

Related to the neutral temperature is the **inversion temperature** θ_i, defined as *the hot-junction temperature at which the EMF returns to the same value it had at the cold junction,* namely zero (assuming the cold junction is at 0°C). Setting $E = 0$ in the first form gives $\theta(a + b\theta) = 0$, whose non-zero root is:

$$\theta = -\frac{a}{b} = 2\theta_N$$

When the cold junction is held at a non-zero temperature θ_{cold}, the EMF curve shifts horizontally but remains symmetric about the neutral temperature, and the inversion temperature becomes:

$$\boldsymbol{\theta_i = 2\theta_N - \theta_{cold}}$$

The Seebeck EMF, its neutral temperature, and its inversion temperature are summarised in the characteristic parabolic graph.

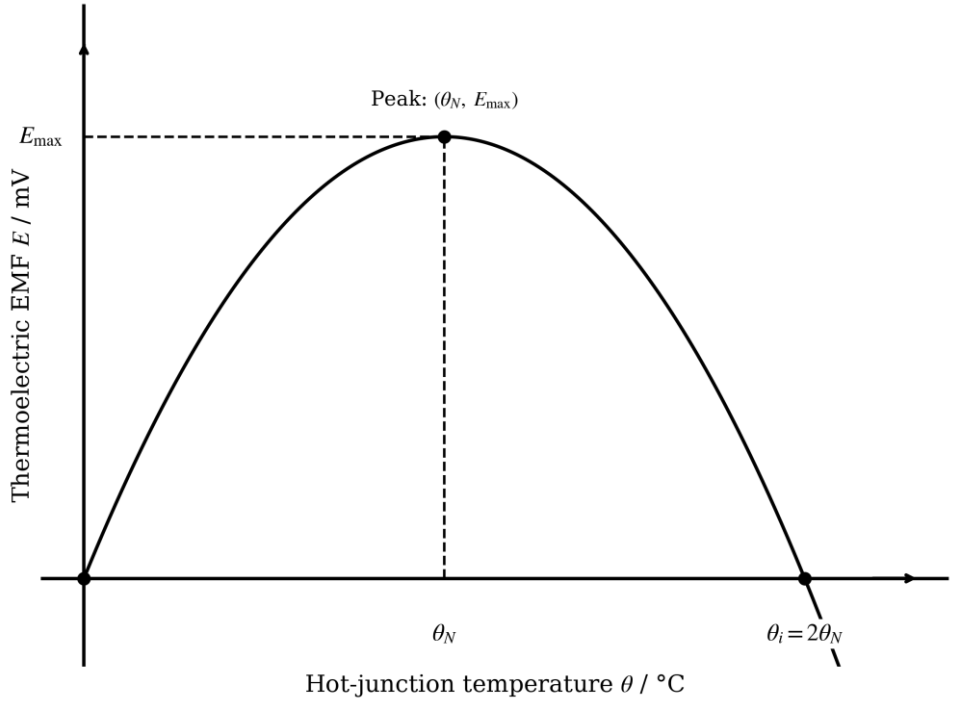


Figure: *The thermoelectric EMF E as a function of hot-junction temperature θ, for a thermocouple whose cold junction is held at $0°C$. The EMF rises from zero at the cold-junction temperature, reaches a maximum E_{max} at the neutral temperature θ_N, and returns to zero at the inversion temperature $\theta_i = 2\theta_N$. Beyond θ_i, the EMF reverses sign, and the thermocouple is then said to be inverted.*

The thermocouple's strengths are its *small size*, its *fast response* (because the junction has almost no thermal mass), its *wide range* (from $-200°C$ to over $1500°C$ depending on the metal pair), and its *ability to measure the temperature at a single point* rather than over a volume. Its weaknesses are its relatively *low accuracy* (typically one or two degrees, not thousandths), its *non-linearity* (forcing the parabolic formula above), and its *need for a stable cold-junction reference*. Thermocouples are the instrument of choice in industrial applications where speed, ruggedness, and high temperature matter more than ultimate precision.

5. Pyrometers

For temperatures above roughly $1500°C$, where even the best thermocouple wires begin to melt or drift, thermometry moves to a different principle altogether. A **pyrometer** *measures temperature by detecting the radiation emitted by a hot body, without making any physical contact with the body itself. The* hottest furnaces, the interiors of blast furnaces, and the surface temperature of the sun itself are measured with pyrometers. The physics behind them is **Wien's displacement law**, which relates the wavelength of peak thermal emission to the absolute temperature, and **Stefan's law,** which relates the total power radiated to the fourth power of the absolute temperature. Both laws are derived formally in Chapter 3; here we simply note that the colour and brightness of a glowing body encode its temperature, and that an optical pyrometer reads that code. Pyrometers are expensive, specialist instruments used chiefly in metallurgy and astrophysics.

Comparing the thermometers at a glance

The five thermometer types we have met differ in range, sensitivity, response time, and intended use. A summary table is the fastest way to see all of them side by side.

Thermometer	Thermometric property	Approximate range	Accuracy	Response time	Typical use
Mercury-in-glass	Volume (column length) of mercury	−39°C to 357°C	±0.1°C	Seconds	Laboratory and clinical work
Alcohol-in-glass	Volume (column length) of alcohol	−115°C to 78°C	±0.5°C	Seconds	Low-temperature and domestic use
Constant-volume gas	Pressure of a fixed mass of gas	−270°C to 1500°C	±0.001°C	Minutes	Metrological standard
Platinum resistance	Electrical resistance of platinum	−200°C to 1000°C	±0.01°C	Seconds	Sub-standard calibration; precision laboratory use
Thermocouple	Seebeck EMF between two metals	−200°C to 1500°C	±1°C	Milliseconds	Industrial use; rapid or point measurements
Optical pyrometer	Wavelength or brightness of emitted radiation	500°C to 3000°C+	±5°C	Instant	Furnaces, molten metal, astrophysics

The table above is not a memorisation exercise. Every entry in it can be recovered from the physics of the thermometer it describes. A thermocouple responds in milliseconds because a junction has almost no thermal mass to warm up. A gas thermometer needs minutes because a gas bulb is a reservoir of molecules that must all reach equilibrium. An alcohol thermometer cannot reach 78°C because ethanol boils there, and the column simply disappears. When in doubt, think about the physics, and the numbers in the table will fall out naturally.

Theory and you have built thermometers together for several pages now. Time for thermal equilibrium with five worked examples to see the instruments in action.

BINDER Example 6

A constant-volume gas thermometer has a gas pressure of 0.900×10^5Pa at the ice point and 1.230×10^5Pa at the steam point. When the bulb is immersed in a water bath, the gas pressure is 1.073×10^5Pa. Calculate the temperature of the bath on the gas scale.

Solution

Substituting $P_0 = 0.900 \times 10^5$Pa, $P_{100} = 1.230 \times 10^5$Pa, and $P_\theta = 1.073 \times 10^5$Pa into the gas-thermometer formula:

$$\theta_{gas} = \left(\frac{P_\theta - P_0}{P_{100} - P_0}\right) \times 100°C$$

$$\theta_{gas} = \left(\frac{1.073 \times 10^5 - 0.900 \times 10^5}{1.230 \times 10^5 - 0.900 \times 10^5}\right) \times 100°C$$

$$\theta_{gas} = 52.4°C$$

The temperature of the bath is $52.4°C$.

Making Sense of the Answer: *The factor of* 10^5 *cancels from top and bottom of the fraction, so the pressures could equally well have been given in any other consistent units, and the answer would have been the same. That is the quiet freedom of interpolation formulas: they care only about the ratio of differences, not about the units in which those differences are measured.*

Think Like a Physicist: *A gas-thermometer calculation is always the same three-step operation: identify the three pressures (at the ice point, at the steam point, and at the unknown temperature), substitute, and compute. The answer will always be a number between the two fixed-point temperatures (assuming the unknown lies between them), so if the calculation produces a value below* $0°C$ *or above* $100°C$*, check whether the unknown pressure really lies in the expected range.*

BINDER Example 7

A platinum resistance thermometer is calibrated at the ice point ($0°C$) and the steam point ($100°C$), with measured resistances of $R_0 = 50.00\Omega$ and $R_{100} = 69.50\Omega$. The variation of platinum resistance with true temperature is known to follow $R_\theta = R_0(1 + a\theta + b\theta^2)$ with $a = 3.98 \times 10^{-3}K^{-1}$ and $b = -5.87 \times 10^{-7}K^{-2}$. When the thermometer is immersed in a liquid whose true temperature (as measured by a gas thermometer) is $50°C$, calculate:

(a) the actual resistance of the platinum coil at this temperature, and

(b) the temperature the platinum thermometer reports on its own empirical scale.

Solution

(a) Substituting $\theta = 50°\text{C}$ into $R_\theta = R_0(1 + a\theta + b\theta^2)$:

$$R_{50} = 50.00\,(1 + (3.98 \times 10^{-3})(50) + (-5.87 \times 10^{-7})(50)^2)\Omega = 59.88\Omega$$

The actual resistance of the platinum coil is 59.88Ω.

(b) Substituting into the interpolation formula for the platinum scale:

$$\theta_R = \left(\frac{R_\theta - R_0}{R_{100} - R_0}\right) \times 100°\text{C} = \left(\frac{59.88 - 50.00}{69.50 - 50.00}\right) \times 100°\text{C}$$

$$\theta_R = \left(\frac{9.88}{19.50}\right) \times 100°\text{C} = 50.67°\text{C}$$

The platinum scale reads $50.67°\text{C}$ when the true temperature is $50°\text{C}$.

Making Sense of the Answer: *The platinum thermometer over-reads by about two-thirds of a degree at the mid-point of its calibration range. This is exactly the kind of disagreement we studied in the previous section, and the numerical size of the disagreement (less than* $1°\text{C}$ *out of* $100°\text{C}$*) confirms that platinum resistance is only mildly non-linear. A mercury thermometer in the same bath would under-read by a similar amount in the opposite direction, and a gas thermometer would sit very close to the true value.*

Think Like a Physicist: *When a problem gives a functional form such as* $R_\theta = R_0(1 + a\theta + b\theta^2)$*, resist the temptation to use it as the thermometer's scale. That functional form tells you the true physical resistance at the true temperature, which is a different question from the reading the thermometer produces. The reading is always obtained by substituting into the linear interpolation formula. The two give different answers precisely because the real physics is non-linear.*

REAL Example 8

Kipanga is in the physics laboratory, measuring the temperature of a hot metal block with a thermocouple. The thermocouple has $a = 40\mu\text{V/°C}$ and $b = -0.04\mu\text{V/°C}^2$ in the form $E = a\theta + b\theta^2$, with the cold junction held in an ice bath at $0°\text{C}$. Kipanga reports a hot-junction EMF of 9.90 mV and proudly announces that the block is at about $470°\text{C}$. Kipute, watching from the next bench, puts down her own work and frowns.

Kipute: *"Kipanga, before you write that number down, do you know what the neutral temperature of that thermocouple is?"*

Kipanga: *"It is where the EMF is maximum. I remember that. Why?"*

Kipute: *"Compute it from the values of* a *and* b *that are stamped on the side of your apparatus."*

Kipanga computes the neutral temperature and finds that it is $500°\text{C}$, only $30°\text{C}$ above his answer. Kipute then asks him to consider three specific problems with his measurement. What was Kipute's argument? Address:

(a) why a reading taken near the neutral temperature is ambiguous,

(b) why the sensitivity of the thermocouple near the neutral temperature is poor, and

(c) why, in view of these two issues, it may be wiser to use a different thermometer entirely for this particular measurement.

Solution

(a) Near the neutral temperature, the EMF-versus-temperature curve is flat (because $dE/d\theta = 0$ at θ_N), which means a small change in the hot-junction temperature produces only a **tiny** change in the EMF. More dangerously, the same EMF corresponds to **two** different hot-junction temperatures, one below θ_N and one above. A reading of 9.90 mV could equally well represent a hot junction at $470°\text{C}$ or one at $530°\text{C}$ (the two temperatures that are symmetric about the neutral temperature $500°\text{C}$), and the thermocouple alone cannot tell Kipanga which is the correct one.

(b) The sensitivity of a thermocouple is $dE/d\theta$, which is the change in EMF per change in temperature. At the neutral temperature, $dE/d\theta = 0$, which means the thermocouple is completely unable to distinguish temperatures at that point. Slightly off the neutral temperature, $dE/d\theta$ is small but non-zero, so a small temperature change produces only a small EMF change. Kipanga's reading lies in a region where even a very careful EMF measurement corresponds to a wide range of possible temperatures, and the instrument's precision is badly degraded.

(c) Given both problems, the thermocouple is the wrong tool for this measurement. A platinum resistance thermometer would have no ambiguity (its resistance rises monotonically with temperature over this range) and excellent sensitivity (of order a tenth of an ohm per degree). If the block is hot enough that a platinum coil cannot be used safely, an optical pyrometer, which reads colour temperature directly and has no neutral-temperature problem at all, is the obvious alternative. Kipanga's confident announcement of $470°\text{C}$ is therefore based on a measurement the thermocouple cannot actually make with the precision he claims, and the sensible response is to repeat the measurement with a different instrument.

Making Sense of the Answer: *Kipute's three points are not abstract worries. Each one is a different failure mode of the same underlying physics, namely that the Seebeck EMF is a parabolic function of temperature and therefore has a peak.*

Below the peak, the EMF rises and the thermometer works; above the peak, the EMF falls and the thermometer lies about the direction of change; near the peak, the EMF is almost constant and the thermometer lies about the size of the change. The right response to all three failures is the same: pick an instrument whose thermometric property is monotonic in the temperature range of interest.

Think Like a Physicist: *Any thermometric instrument has a working range, and the working range ends well before the property stops changing altogether. For a thermocouple, the working range ends some distance below the neutral temperature, precisely to keep the sensitivity* $\mathrm{dE/d\theta}$ *comfortably above zero. Whenever you are asked to use a thermocouple near its neutral temperature, the answer is almost always that you should not. A physicist who knows the shape of a curve knows which parts of the curve are safe to measure on, and which parts are not.*

HOT Example 9

A thermocouple has its cold junction held in an ice-water bath at 0°C and its thermoelectric EMF follows $\mathrm{E = a\theta + \frac{1}{2}b\theta^2}$, with $\mathrm{a = 52.0\mu V/°C}$ and $\mathrm{b = -0.040\mu V/°C^2}$. Calculate (a) the neutral temperature, (b) the maximum EMF, and (c) the inversion temperature.

Solution

(a) The neutral temperature is where $\frac{\mathrm{dE}}{\mathrm{d\theta}} = 0$. Differentiating the given form of E:

$$\frac{\mathrm{dE}}{\mathrm{d\theta}} = \mathrm{a + b\theta}$$

Setting $\frac{\mathrm{dE}}{\mathrm{d\theta}} = 0$ at $\theta = \theta_N$:

$$\mathrm{a + b\theta_N = 0 \quad \Rightarrow \quad \theta_N = -\frac{a}{b}}$$

$$\theta_N = -\frac{52.0\mu V/°C}{-0.040\mu V/°C^2} = 1300°C$$

The neutral temperature is $1300°C$.

(b) The maximum EMF is E evaluated at $\theta = \theta_N$. Substituting $\theta_N = -\frac{a}{b}$ back into the given form:

$$E_{max} = a\left(-\frac{a}{b}\right) + \frac{1}{2}b\left(-\frac{a}{b}\right)^2 = -\frac{a^2}{b} + \frac{a^2}{2b} = -\frac{a^2}{2b}$$

Substituting numerical values:

$$E_{max} = -\frac{(52.0\mu V/°C)^2}{2(-0.040\mu V/°C^2)} = 3.38 \times 10^4 \mu V = 33.8mV$$

The maximum EMF is 33.8mV.

(c) The inversion temperature is where the EMF returns to 0. With the cold junction at $0°C$, setting $E = 0$:

$$\mathrm{a\theta + \frac{1}{2}b\theta^2 = 0 \quad \Rightarrow \quad \theta\left(a + \frac{1}{2}b\theta\right) = 0}$$

The non-zero root is:

$$\theta_i = -\frac{2a}{b} = 2\theta_N$$

$$\theta_i = 2(1300°C) = 2600°C$$

The inversion temperature is $2600°C$.

Making Sense of the Answer: *A neutral temperature of* $1300°C$ *is well above the melting point of most metals a Form 5 student has met, which is exactly where the thermocouple's usefulness lies. The inversion temperature at* $2600°C$ *is higher than any practical thermocouple can actually survive, because the metal wires themselves would have melted long before the hot junction reaches that temperature. The inversion temperature is therefore an idealised mathematical root of the parabolic fit, not an operational limit; the real working limit is set by whichever wire melts first.*

Think Like a Physicist: *When a problem gives you* a *and* b *together with a form of* E*, the first move is always to identify which form (with or without the factor of* $1/2$*) the constants are fitted to. The same pair of numerical values of* a *and* b *produces answers that differ by a factor of two between the two forms, and mixing forms is one of the commonest mistakes in thermocouple problems. Read the formula given in the question before reaching for any memorised shortcut.*

HOT Example 10

At an unknown temperature, a constant-volume gas thermometer gives a pressure reading of $P_\theta = 1.120 \times 10^5 Pa$ (where the ice and steam point pressures are $P_0 = 0.800 \times 10^5 Pa$ and $P_{100} = 1.200 \times 10^5 Pa$), and a platinum resistance thermometer in the same bath reads $R_\theta = 58.50\Omega$ (where $R_0 = 50.00\Omega$ and $R_{100} = 70.00\Omega$).

(a) Calculate the temperature of the bath on the gas scale and on the platinum resistance scale.

(b) Using the gas-scale value as the true temperature, calculate the percentage discrepancy of the platinum-scale reading.

(c) Assess the reliability of the platinum resistance thermometer reading and state what action a physicist should take.

Solution

(a) Substituting into the interpolation formula for each thermometer.

Gas scale:

$$\theta_{\text{gas}} = \left(\frac{P_\theta - P_0}{P_{100} - P_0}\right) \times 100^\circ\text{C} = \left(\frac{1.120 - 0.800}{1.200 - 0.800}\right) \times 100^\circ\text{C} = 80.0^\circ\text{C}$$

The temperature of the bath on the gas scale is 80.0°C.

Platinum resistance scale:

$$\theta_{\text{R}} = \left(\frac{R_\theta - R_0}{R_{100} - R_0}\right) \times 100^\circ\text{C} = \left(\frac{58.50 - 50.00}{70.00 - 50.00}\right) \times 100^\circ\text{C} = 42.5^\circ\text{C}$$

The temperature of the bath on the platinum resistance scale is 42.5°C.

(b) Percentage discrepancy of the platinum reading relative to the true gas-scale value:

$$\text{Discrepancy} = \frac{\theta_{\text{R}} - \theta_{\text{gas}}}{\theta_{\text{gas}}} \times 100\% = \frac{42.5 - 80.0}{80.0} \times 100\% = -46.9\%$$

The percentage discrepancy is 46.9% below the true temperature.

(c) The platinum-scale reading is 46.9% below the true temperature. A discrepancy of this size is far too large to result from ordinary non-linearity of platinum resistance, which for a properly constructed coil would be less than a few percent over this range. The most likely explanation is that the platinum coil has been damaged, contaminated, or badly calibrated. A physicist in this situation would send the platinum resistance thermometer for recalibration or replacement.

Making Sense of the Answer: *The numbers in this problem are chosen to make the two thermometers disagree visibly. In a real laboratory, a* 50% *disagreement between a gas thermometer and a platinum resistance thermometer would not be a textbook illustration of non-linearity; it would be a sign that something is wrong with the instrument. Real platinum resistance thermometers agree with gas thermometers to better than a degree over their working range, and the small difference between the two can be calibrated out for high-precision work.*

Think Like a Physicist: *When two independent thermometers disagree by a few tenths of a degree, suspect their physics (non-linearity of the properties). When they disagree by several degrees, suspect their calibration. When they disagree by tens of degrees or more, suspect the instrument itself. Different sizes of disagreement point to different causes, and a physicist's first act on meeting a large discrepancy is not to invoke subtle theory but to check whether the apparatus is actually working.*

We have now met every instrument a student is likely to encounter in a laboratory or on an examination paper. Each one uses a different thermometric property, each one has its own range and its own non-linearity, and each one agrees with the others only at the fixed points on which it was calibrated. The disagreement between them is not a flaw of the thermometers, it is a flaw of the very idea of calibrating a scale on two points and assuming linearity in between. If we want a single temperature scale that every thermometer can be made to agree with, a scale that does not depend on the liquid inside the glass or the metal inside the wire, we have to escape the empirical framework altogether. The next section shows how physicists did exactly that, by choosing a single fixed point defined by the physics of matter itself, and by building a scale on it that is universal across every substance in nature.

THE THERMODYNAMIC SCALE AND THE TRIPLE POINT

Every empirical thermometer we have met so far is a good instrument with a small private flaw. The flaw is that it reports its own physics. A mercury thermometer reports how mercury expands, a platinum thermometer reports how platinum's resistance changes, a gas thermometer reports how a fixed mass of gas responds. Calibrating any of these instruments on two fixed points produces a scale that agrees with the other scales only at those two points. Between the fixed points, each thermometer tells a slightly different story, and nothing in the empirical approach can say which story is the right one.

This is not a small inconvenience. A temperature of "50°C on the mercury scale" is a legitimate quantity, but so is "50°C on the platinum scale", and the two numbers refer to physically different states. If the world contained only one thermometric substance we would be fine; since it does not, we are forced to ask a harder question. *Is there a temperature scale that does not depend on the thermometric substance at all? A scale that every properly constructed thermometer would converge towards in the limit of perfect construction?* This section answers that question with a clear yes, and builds the scale explicitly.

The problem with two fixed points

The empirical Celsius scale uses two fixed points, the ice point and the steam point. Both are sensitive to external conditions in ways that the casual reader does not always appreciate. The **ice point** (*the temperature at which pure water and pure ice coexist at equilibrium*) depends on the applied pressure, because changing the pressure on ice shifts the melting temperature slightly. The **steam point** (*the temperature at which pure water boils*) depends very strongly on the applied pressure, because **boiling** *is the point where the vapour pressure of water equals the atmospheric pressure above it.* A laboratory 2000m above

sea level and a laboratory at sea level will therefore measure two slightly different ice points and two significantly different steam points, unless both are carefully corrected to standard atmospheric pressure.

Even worse, the "ice point" as normally prepared in a laboratory contains dissolved air in the water, which depresses the freezing temperature by about 0.01°C relative to pure water with no air dissolved in it. This is small, but it is also uncorrectable without careful preparation, and a scale that depends on how carefully the ice-water mixture was prepared is a scale with an unwelcome human fingerprint on it.

The thermodynamic scale escapes all of this by refusing to use either the ice point or the steam point as a reference. Instead, it uses a single fixed point defined by a physical state that is unique in nature, and from that single point it builds the entire scale.

The triple point of water

A **phase diagram** for water plots pressure on the vertical axis and temperature on the horizontal axis, and marks the curves along which two different phases of water coexist in equilibrium. Along the fusion curve, solid and liquid coexist; along the vapourisation curve, liquid and vapour coexist; along the sublimation curve, solid and vapour coexist. The three curves are not independent: they must meet at a single point, and at that point all three phases of water (ice, liquid water, and water vapour) coexist simultaneously in equilibrium. This single point is called the **triple point of water**.

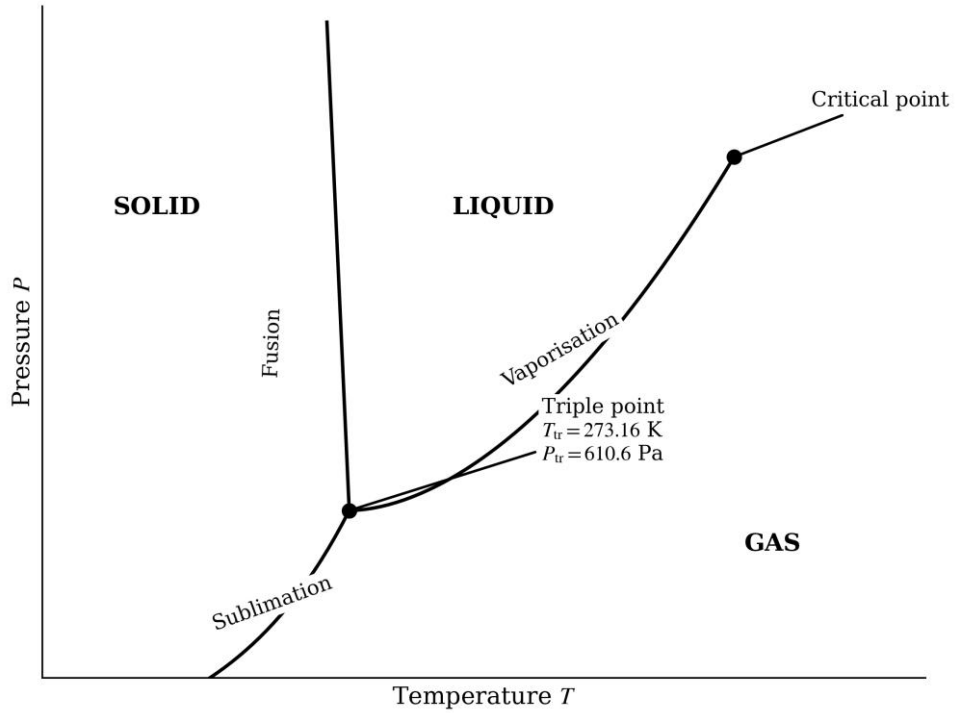


Figure: ***The phase diagram of water.*** *The three phase-boundary curves (fusion, vapourisation, sublimation) meet at the triple point, at which solid, liquid, and vapour coexist in equilibrium at a unique temperature and pressure. The vapourisation curve ends at the critical point, beyond which liquid and vapour become indistinguishable.*

The triple point of water occurs at a unique combination of temperature and pressure, specifically at $T_{tr} = 273.16$K and $P_{tr} \approx 610.6$Pa (about 0.006atm). These two numbers are not adjustable. They are properties of water itself, determined by the intermolecular forces in the water molecule, and any laboratory anywhere in the universe that correctly prepares a triple-point cell (a sealed vessel containing pure water with ice, liquid, and vapour in equilibrium) will measure the same temperature and the same pressure to within the precision of its instruments.

Compare this with the ice point. An ice-point bath depends on the purity of the water, the presence of dissolved air, and the applied atmospheric pressure, and each of these introduces a small correction that must be handled carefully. A triple-point cell depends on none of them. Once the cell is sealed, the pressure is fixed by the vapour pressure of water in contact with ice, the temperature is fixed by the simultaneous coexistence of all three phases, and nothing outside the cell can change either value without breaking the triple-point condition. The triple point is therefore *reproducible to an unprecedented precision*, and it is the only fixed point used to define the modern thermodynamic scale.

Why the ice point sits at 273.15 K, not 273.16 K

A small numerical detail is worth clearing up before we go further. The triple point of water is at exactly 273.16K, by definition of the kelvin. The ice point of pure water at standard atmospheric pressure sits at 273.15K, which is exactly 0.01K (i.e. 0.01°C) below the triple point. The 0.01K difference comes from two small corrections.

First, the ice point uses pure water and ice at standard atmospheric pressure (1.013×10^5 Pa), not at the triple-point pressure of 610.6Pa, and the higher pressure lowers the freezing temperature slightly.

Second, the ice point uses water that has been in contact with air, and the dissolved air lowers the freezing point by a further small amount.

Adding the two corrections gives approximately the 0.01K gap between the two reference temperatures.

The thermodynamic-scale formula

The thermodynamic scale is built on the observation that the pressure of an ideal gas held at constant volume is directly proportional to its absolute temperature. For an ideal gas at constant volume, Gay-Lussac's law gives

$$P \propto T \quad \Rightarrow \quad \frac{P}{T} = \text{constant}$$

This relation becomes exact in the limit of vanishing gas density and is very nearly exact for any dilute real gas. Now imagine a constant-volume gas thermometer whose bulb is first immersed in a triple-point cell and gives a pressure reading $\mathbf{P_{tr}}$, and then immersed in a body at an unknown temperature $\mathbf{T}$ where it gives a pressure reading $\mathbf{P_\theta}$. Applying the constant-ratio relation at both temperatures:

$$\frac{P_\theta}{T} = \frac{P_{tr}}{T_{tr}}$$

Solving for T:

$$T = T_{tr} \times \frac{P_\theta}{P_{tr}}$$

Substituting the defined value $T_{tr} = 273.16\text{K}$:

$$\mathbf{T = 273.16K \times \frac{P_\theta}{P_{tr}}}$$

This single equation defines the thermodynamic temperature T of any body, using only the pressure reading of an ideal gas thermometer at two states: the state of the body, and the triple point of water.

Notice that the whole scale has been built from a single fixed point. There is no ice point, no steam point, no interpolation. The scale is absolute, in the sense that $T = 0$ corresponds to $P_\theta = 0$, which is the state an ideal gas would reach if its pressure could be brought to zero at constant volume. This temperature is called **absolute zero**, and it is the lowest temperature physically permitted by the laws of thermodynamics.

Absolute zero and the Celsius–Kelvin conversion

Absolute zero is not a temperature that any real gas thermometer can actually reach, because *every real gas **liquefies** or **solidifies** long before its pressure approaches zero.* But by measuring the pressure of a gas at several high temperatures and extrapolating the straight line $P \propto T$ backwards to $P = 0$, the horizontal-axis intercept gives the temperature of absolute zero. The result, common to every real gas in the dilute limit, is $T = 0\text{K}$, which on the Celsius scale corresponds to $\theta = -273.15°\text{C}$.

The Celsius scale and the thermodynamic (kelvin) scale are therefore related by a simple additive shift:

$$\mathbf{T(K) = \theta(°C) + 273.15}$$

A change of 1K and a change of 1°C are numerically equal, because the two scales have the same unit size; only their zero points differ. The ice point is 273.15K on the absolute scale and 0°C on the Celsius scale. The steam point is 373.15K and 100°C. The triple point is 273.16K and 0.01°C. Absolute zero is 0K and −273.15°C. These correspondences, together with the older Fahrenheit scale still used in some countries, are summarised in the following diagram.

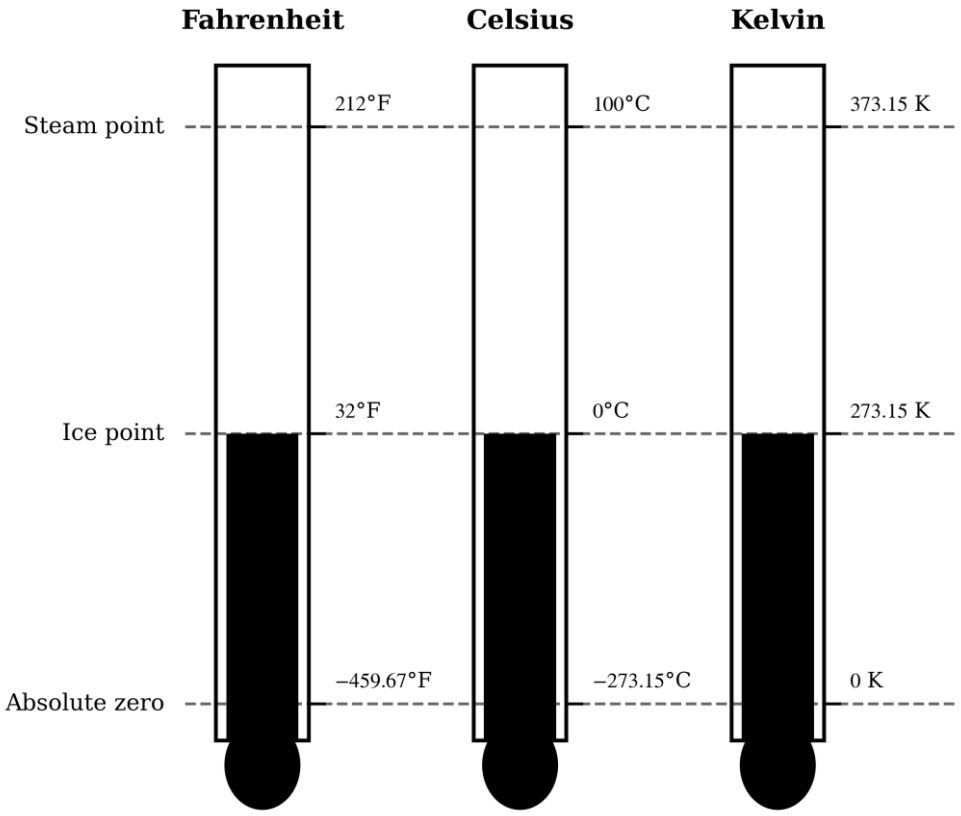


Figure: ***The three temperature scales side by side.*** *The Fahrenheit scale places the ice point at* $32^\circ F$ *and the steam point at* $212^\circ F$*. The Celsius scale uses* $0^\circ C$ *and* $100^\circ C$*. The thermodynamic (Kelvin) scale uses* $273.15K$ *and* $373.15K$*, with absolute zero as its true zero. Notice that absolute zero is not labelled on the Celsius or Fahrenheit scales as a defining point; it is merely a value those scales happen to reach.*

Theory and you have travelled from empirical scales to a substance-independent absolute scale in a handful of pages. Time for thermal equilibrium with two worked examples to settle the ideas.

BINDER Example 11

The modern thermodynamic scale uses only the triple point of water as its fixed reference, while the older Celsius scale used two fixed points, the ice point and the steam point.

(a) Explain why the triple point is a more reliable fixed reference than the ice point.

(b) Show, using the constant-volume gas thermometer equation, how a single fixed point combined with the ideal-gas relation is sufficient to define an entire temperature scale.

Solution

(a) The ice point depends on three external conditions that are difficult to control. First, the temperature at which pure water and pure ice coexist depends on the applied pressure, and any deviation from standard atmospheric pressure shifts it slightly. Second, water in contact with air absorbs dissolved oxygen and nitrogen, which depress the freezing point by roughly $0.01^\circ C$. Third, the purity of the water itself affects the freezing temperature. A laboratory that wishes to use the ice point as a reference must correct for all three, and the corrections introduce small laboratory-to-laboratory variations.

The triple point, by contrast, depends on none of these. At the triple point, pure water, pure ice, and water vapour coexist in a single sealed vessel at a temperature and pressure that are both determined uniquely by the physics of the water molecule itself. The vessel is sealed, so no air can dissolve in the water. The pressure inside the cell is fixed at $610.6Pa$ by the vapour pressure of water in contact with ice, so no external atmosphere can disturb it. The temperature inside the cell is fixed at $273.16K$ by the requirement that all three phases coexist in equilibrium. Any laboratory that prepares a triple-point cell correctly will measure the same temperature and the same pressure, regardless of altitude, weather, or local atmospheric composition. This is what makes the triple point "fully reproducible" in the physicist's sense: the same physical state gives the same number everywhere.

(b) The thermodynamic scale is built on the constant-volume gas relation $P \propto T$, which for an ideal gas gives $\frac{P}{T} = \text{constant}$. Applying this at the triple point and at an unknown temperature T:

$$\frac{P_{tr}}{T_{tr}} = \frac{P_\theta}{T} \quad \Rightarrow \quad T = T_{tr} \times \frac{P_\theta}{P_{tr}}$$

With the defined value $T_{tr} = 273.16\text{K}$:

$$T = 273.16\text{K} \times \frac{P_\theta}{P_{tr}}$$

One fixed point (the triple point, which fixes T_{tr}) and one physical law (the ideal gas relation) together determine T at every other pressure. No second fixed point is needed, because the ideal gas law itself fixes the shape of the P-versus-T relation. The empirical Celsius scale needed two fixed points because it made no assumption about the shape of its thermometric property's variation with temperature; it had to calibrate that shape at two points and interpolate linearly. The thermodynamic scale needs only one fixed point because the ideal-gas law tells us in advance that the shape is linear through the origin.

Making Sense of the Answer: *The move from two fixed points to one is not just a cosmetic simplification. It is the move from an empirical to a fundamental temperature scale, because the ideal-gas relation is a law of physics and not a property of a particular substance. Every thermometer ever built, if used correctly, will converge on the thermodynamic temperature in the limit of vanishing gas density. The triple point is simply the anchor the world has agreed to peg the scale to, and* $273.16K$ *is the number that makes the new scale numerically consistent with the older Celsius scale at the ice point.*

Think Like a Physicist: *When you are told that a scale has been defined by a single reference point and a physical law, look for the law carefully. The law is doing half the work. In the thermodynamic scale, the law is Gay-Lussac's* $P \propto T$*; in other contexts it might be* $E = mc^2$ *or* $\lambda T = \text{constant}$*. The reference point alone is never enough to build a scale; it is the combination of a reference point and a law that makes the construction possible.*

BINDER Example 12

A constant-volume gas thermometer has a pressure of $1.500 \times 10^4\text{Pa}$ at the triple point of water. The bulb is then immersed in a pan of boiling water at an unknown temperature, and the pressure rises to $2.050 \times 10^4\text{Pa}$.

(a) Find this temperature on the Kelvin scale.

(b) Convert the answer to the Celsius scale.

(c) Comment on whether the result is consistent with the boiling point of water at standard atmospheric pressure.

Solution

(a) Substituting into the thermodynamic-scale formula $T = 273.16\text{K} \times \frac{P_\theta}{P_{tr}}$:

$$T = 273.16\text{K} \times \frac{2.050 \times 10^4\text{Pa}}{1.500 \times 10^4\text{Pa}} = 373.32\text{K}$$

The temperature on the Kelvin scale is 373.3K.

(b) Converting to Celsius:

$$\theta = T - 273.15 = 373.32 - 273.15 = 100.17°\text{C} \approx 100.2°\text{C}$$

The temperature on the Celsius scale is $100.2°\text{C}$.

(c) The boiling point of water at standard atmospheric pressure is $100°\text{C}$, and the thermometer reads $100.2°\text{C}$. The two values agree to within $0.2°\text{C}$, which is consistent with the thermometer measuring a body of boiling water at close to standard pressure. The small positive discrepancy could come from a slightly elevated atmospheric pressure in the laboratory (which would raise the boiling point above $100°\text{C}$), or from a small non-ideal-gas correction in the thermometer, or from ordinary experimental error. In any case, the reading is physically sensible for boiling water.

Making Sense of the Answer: *The answer matches a temperature the student already recognises, which is the whole purpose of the exercise. The thermodynamic formula is not a mystery: it is the ideal-gas relation rearranged to give a temperature from a pressure ratio. When the ratio is roughly 1.37, the temperature is roughly 1.37 times the triple-point temperature, which puts it close to the boiling point of water. Whenever a calculation produces a number close to a physically meaningful temperature like 273K or 373K, pause and check which physical state the body is probably in.*

Think Like a Physicist: *A gas thermometer calculation on the thermodynamic scale is always the same two-step operation: take the ratio of the two pressures, and multiply by* 273.16K*. The whole calculation is one multiplication, yet it is one of the most powerful measurements in physics, because it measures an absolute temperature without reference to any other substance. If a problem gives you a pressure ratio and asks for a thermodynamic temperature, do not reach for the Celsius interpolation formula; the answer is always* 273.16K *times the ratio.*

We have now escaped the empirical framework and reached a temperature scale that every laboratory in the world can measure against the same reference. But the thermodynamic scale is not the end of the story. In practice, physicists still use mercury thermometers, platinum coils, and thermocouples, all of which are empirical instruments. *When exactly do these agree with*

the thermodynamic scale, and when do they diverge? The next section takes the four thermometers we have built, sets them side by side at the same temperature, and works out how their readings compare.

COMPARISON AND LIMITATIONS OF THERMOMETRIC SCALES

We now have a zoo of thermometric scales at our disposal. Every empirical thermometer from the previous sections gives its own scale, and alongside them we have the thermodynamic (Kelvin) scale, which was derived from first principles and depends on no particular substance. The empirical scales agree with each other, and with the thermodynamic scale, only at the fixed points on which they were calibrated. Everywhere else, they disagree by small but systematic amounts. This section examines the pattern of agreement and disagreement, identifies which thermometer sits closest to the thermodynamic truth, and states the limitations that remain even for the best empirical instrument.

Where the scales agree and where they diverge

Consider four thermometers: a mercury-in-glass thermometer, a platinum resistance thermometer, a constant-volume gas thermometer, and a thermocouple. Suppose all four are carefully calibrated at the ice point and the steam point, and all four are then immersed in a water bath at some intermediate temperature. What readings do they give?

At the ice point and the steam point, all four thermometers give the same reading by construction. This is trivial, because calibration forces every thermometer to agree at its two calibration temperatures. The interesting question is what happens between them.

At every intermediate temperature, the four thermometers give four slightly different readings. The mercury thermometer reads slightly below the true temperature (because the **expansion of mercury is sub-linear** in temperature between the fixed points), the platinum resistance thermometer reads slightly above (because the **resistance of platinum is super-linear),** the thermocouple reads an amount that depends on where the measurement falls on its parabolic EMF curve, and the gas thermometer reads very nearly correctly, closer to the truth than any of the other three. The pattern is not random. Each thermometer's deviation reflects the shape of its own thermometric property, and *the size of the deviation reflects how non-linear that property is.*

Why the gas thermometer is closest to the truth

Among all empirical thermometers, the constant-volume gas thermometer enjoys a special status because its thermometric property (the pressure of a fixed mass of gas at constant volume) is almost exactly linear in temperature. For an ideal gas, the relation $P \propto T$ is exact, and the gas-scale temperature equals the thermodynamic temperature identically. Real gases deviate slightly from ideal behaviour because their molecules attract each other weakly, and this causes the gas-thermometer scale to differ from the thermodynamic scale by a small correction.

The correction shrinks as the amount of gas in the bulb is reduced, because in the low-density limit the gas molecules are so far apart that their interactions become negligible and the ideal-gas law becomes exact. In the limit as the gas quantity tends to zero, the gas thermometer and the thermodynamic scale become identical. This is why the thermodynamic scale is said to be **defined** by the ideal-gas relation taken in the low-density limit, and it is why the gas thermometer is used as the primary standard against which all other empirical thermometers are calibrated.

Limitations that remain

Even for the gas thermometer, three limitations remain.

1) The low-density limit can only be approached, never reached, because a thermometer with zero gas in its bulb would produce no pressure and no reading. In practice, multiple measurements are taken at successively lower gas densities and the results are extrapolated to zero density, which introduces a small but finite error.

2) The gas thermometer is slow, delicate, and impractical for routine use, which means most laboratory and industrial measurements use faster thermometers (mercury, platinum, thermocouple) that are in turn calibrated against the gas thermometer.

3) At very high temperatures (above roughly $1000\,^{\circ}\mathrm{C}$) no gas thermometer can survive, and the thermodynamic scale must be extended using other physical principles, such as the Stefan law of blackbody radiation or the Wien displacement law that we will meet in Chapter 3.

For the other empirical thermometers, the principal limitation is non-linearity. A mercury thermometer calibrated at the ice point and the steam point will read a few tenths of a degree away from the truth at the mid-point of its range, and the error is systematic rather than random. For most laboratory work, this error is small enough to ignore; for precision measurements, a correction table is constructed by comparing the thermometer with a gas thermometer at many intermediate temperatures, and the correction is applied to every reading.

Practical consequences for the student

For a student, two practical consequences follow from the comparison above.

The first is that a question asking for "the temperature" of a body requires the student to identify which scale is being used.

- When the question specifies a thermometer by name, the answer is that thermometer's empirical reading, obtained by substituting into its interpolation formula.
- When the question speaks of "the true temperature" or "the thermodynamic temperature", the answer is the gas-thermometer reading (or equivalently the Kelvin temperature computed from the triple-point formula).
  Treating the two as interchangeable is one of the commonest mistakes in thermometry questions.

The second is that when two thermometers are calibrated at the same fixed points and give different readings at an intermediate temperature, neither of them is broken. The disagreement is a consequence of their different physics, and the only legitimate response is to quote both readings alongside the name of the scale. A single "temperature" without a scale is an incomplete quantity, much like a length without a unit.

Theory and you have surveyed every thermometer in this chapter from a high vantage point. Two worked examples now bring the comparison down to ground level.

BINDER Example 13

A platinum resistance thermometer and a constant-volume gas thermometer are both calibrated at the ice point (0°C) and the steam point (100°C). A student argues: *"Since both thermometers use the same two fixed points, they must agree at every temperature between them."* Explain precisely why this student is wrong, and identify the single condition under which the two scales would agree at every intermediate temperature.

Solution

The student's argument is wrong because calibration at two fixed points forces the two thermometers to agree only at those two points, not everywhere in between. The interpolation formula assumes the thermometric property varies linearly with temperature, but the actual variation is linear for **neither** thermometer exactly. The platinum resistance is slightly super-linear with temperature, so the platinum-scale reading bulges above the straight line, while the gas pressure at constant volume is very nearly linear with temperature and sits very close to the straight line. The two scales therefore disagree at every intermediate temperature, by amounts that reflect the shape of each thermometric property's graph.

The single condition under which the two scales would agree at every temperature is that both thermometric properties must be **exactly linear** functions of the true thermodynamic temperature, with the same slope per unit of calibration range. In practice, this is true only in an idealised limit: if the gas in the gas thermometer were truly ideal (no inter-molecular forces) and the resistance of platinum were truly linear in temperature (no higher-order terms in θ), the two scales would coincide. Real gases approach the ideal limit at low density, but real platinum does not become perfectly linear by any such limit, so the two scales always retain a small non-zero discrepancy.

Making Sense of the Answer: *Two fixed points are the minimum requirement for calibrating a scale, but they are not enough to make two different scales agree everywhere. The formula that interpolates between them is a straight line, and two curved graphs that share the same two endpoints can still differ everywhere in between. The student's mistake is confusing agreement at the endpoints with agreement everywhere, which is a very common error for anyone who has not yet thought carefully about what linear interpolation actually promises.*

Think Like a Physicist: *Whenever two thermometers are calibrated at the same fixed points but show different readings at an intermediate temperature, the disagreement reveals the non-linearity of at least one of the two thermometric properties. This is not a weakness of the instruments; it is information about the physics of the substances inside them. A physicist who sees two thermometers disagree does not declare one of them broken; they record both readings, note the disagreement, and use it to test their understanding of the underlying material properties.*

HOT Example 14

The resistance of a platinum coil at the ice point, at the steam point, and in a certain room is $R_0 = 63.0\Omega$, $R_{100} = 75.0\Omega$, and $R_\theta = 65.09\Omega$ respectively. A constant-volume gas thermometer at the same three temperatures gives pressures $P_0 = 8.00 \times 10^6\text{Pa}$, $P_{100} = 1.10 \times 10^7\text{Pa}$, and $P_\theta = 8.51 \times 10^6\text{Pa}$.

(a) Calculate the room temperature on the platinum resistance scale and on the gas scale.

(b) Explain why the two readings differ.

(c) State which of the two readings is closer to the true thermodynamic temperature and justify your choice.

Solution

(a) Applying the interpolation formula to each thermometer.
Platinum scale:

$$\theta_R = \left(\frac{R_\theta - R_0}{R_{100} - R_0}\right) \times 100°\text{C} = \left(\frac{65.09 - 63.0}{75.0 - 63.0}\right) \times 100°\text{C} = 17.42°\text{C}$$

The room temperature on the platinum resistance scale is $17.42°\text{C}$.
Gas scale:

$$\theta_{\text{gas}} = \left(\frac{P_\theta - P_0}{P_{100} - P_0}\right) \times 100°\text{C} = \left(\frac{8.51 \times 10^6 - 8.00 \times 10^6}{1.10 \times 10^7 - 8.00 \times 10^6}\right) \times 100°\text{C} = 17.00°\text{C}$$

The room temperature on the gas scale is $17.00°\text{C}$.

(b) The two readings differ by $0.42°\text{C}$ because the two thermometric properties are not perfectly linear with true temperature. The resistance of platinum is slightly super-linear between the fixed points, which means the platinum resistance at the true temperature of $17°\text{C}$ is slightly higher than a linear interpolation from the fixed points would predict. Substituting this slightly elevated resistance into the interpolation formula returns a reading of $17.42°\text{C}$, which is above the true value. The gas pressure at constant volume, by contrast, is very nearly linear with true temperature, so the gas reading of $17.00°\text{C}$ sits very close to the truth. The disagreement between the two scales is therefore a genuine consequence of the non-linear physics of platinum, not a calibration error.

(c) The gas-scale reading of $17.00°\text{C}$ is closer to the true thermodynamic temperature. The pressure of a fixed mass of gas at constant volume is nearly proportional to the thermodynamic temperature, with deviations from ideal behaviour that become negligible in the low-density limit. The resistance of platinum, by contrast, has no comparable principle of approach to the thermodynamic scale; its non-linearity is a property of the metal itself and does not disappear with any limiting procedure. The gas thermometer is therefore the instrument of choice when the true thermodynamic temperature is needed, and the platinum reading must be corrected against a gas thermometer before it can be trusted for precision work.

Making Sense of the Answer: *The two readings differ by less than half a degree out of roughly seventeen, which is less than three percent. In a physics laboratory this is a large and noticeable disagreement; in a hospital it would be undetectable; in a kitchen it would be invisible. The size of the disagreement is therefore a guide to which thermometer is appropriate for which task. Precision work demands the gas thermometer or a platinum thermometer carefully corrected against it; routine work is perfectly well served by either.*

Think Like a Physicist: *Whenever a problem gives you readings of two thermometers at the same three states, the calculation is always two parallel interpolation formulas. The interesting physics lies in explaining the difference between the two answers, not in the arithmetic of the two substitutions. Practise looking at the size of the difference and identifying which thermometric property is responsible. A disagreement of a few tenths of a degree almost always reveals a non-linearity, a disagreement of a few degrees reveals a miscalibration, and a disagreement of tens of degrees reveals a broken instrument.*

We have now closed the book on thermometry. Every thermometer we will ever need has been built, every scale we will ever quote has been defined, and every limitation has been made explicit. What remains is to leave temperature behind and move on to the other great thermal quantity, heat itself, and to learn how to compute the heat absorbed or released when a body changes its temperature. The next section introduces the specific heat capacity, which tells us how much heat it takes to warm one kilogram of a substance by one degree, and turns the thermometer from a passive reader of temperature into a tool for quantitative thermal calculation.

SPECIFIC HEAT CAPACITY

We have spent five sections learning how to measure temperature honestly. It is time now to put the thermometer to work, because temperature on its own is only half the thermal story. The other half is heat, the flow of energy from one body to another because of a temperature difference, and the question we have not yet asked is the most practical one of all. *When heat flows into a body, by how much does the body's temperature rise? When heat flows out, by how much does it fall?* The answer depends on the substance, and the quantity that encodes the answer is called the **specific heat capacity**.

Defining specific heat capacity

Experiments show that the heat ΔQ required to raise the temperature of a body is proportional to two quantities: the mass m of the body, and the temperature rise ΔT. The proportionality is written as

$$\Delta Q = mc\Delta T$$

where c is a substance-dependent constant called the specific heat capacity. Rearranging:

$$c = \frac{\Delta Q}{m\Delta T}$$

In words, the **specific heat capacity** *is the heat required to raise the temperature of one kilogram of the substance by one kelvin* (or equivalently one degree Celsius, since the two intervals are numerically equal). Its SI unit is $\text{J kg}^{-1}\text{K}^{-1}$.

The specific heat capacity is a property of the substance, not of the particular sample. Two kilograms of water and two hundred kilograms of water both have the same specific heat capacity of $4200\text{J kg}^{-1}\text{K}^{-1}$; the difference is that warming the larger sample by one degree takes a hundred times more heat, because it has a hundred times more mass. The formula $\Delta Q = mc\Delta T$ handles this automatically: the heat required increases with the mass for a fixed temperature rise, and with the temperature rise for a fixed mass.

Thermal capacity

For a particular sample of matter, it is sometimes convenient to bundle the mass and the specific heat capacity together into a single quantity called the **thermal capacity**:

$$C = mc$$

The **thermal capacity** (unit J K^{-1}) *is the heat required to raise the whole sample's temperature by one kelvin,* and the heat equation then reads $\Delta Q = C\Delta T$. The specific heat capacity belongs to the substance; the thermal capacity belongs to the sample. Two saucepans made of identical cast iron have the same specific heat capacity but different thermal capacities if one is heavier than the other. Thermal capacity is the quantity that actually decides how much heat your kitchen needs to warm something, and it reappears constantly in the worked examples that follow.

A quick note on notation: For solids and liquids, the specific heat capacity is almost independent of how the heating is carried out, whether at constant volume or constant pressure, because solids and liquids barely change their volume with temperature. For gases, however, the two differ noticeably, and physicists distinguish between c_p (specific heat at constant pressure) and c_v (specific heat at constant volume). We will meet this distinction properly in Chapter 4 when we develop the kinetic theory of gases; for the chapter in hand, the word "specific heat capacity" without subscript means the common value used for solids and liquids.

Measuring the specific heat capacity: the method of mixtures

The oldest and most intuitive way to measure the specific heat capacity of a solid is the **method of mixtures**. A known mass of the substance in question is heated to a known high temperature, then dropped into a calorimeter containing water at a known lower temperature, and the final equilibrium temperature is recorded. The principle behind the method is the conservation of energy: in an insulated system, the heat lost by the hot body equals the heat gained by the cold water plus the heat gained by the calorimeter itself.

Let m_1 be the mass of the hot substance, c_1 its (unknown) specific heat capacity, T_1 its initial high temperature, and T_f the final equilibrium temperature. Let m_2 be the mass of the water, $c_2 = 4200\,\text{J kg}^{-1}\text{K}^{-1}$ its specific heat capacity, T_2 its initial temperature (equal to the calorimeter's initial temperature). Let m_{cal} be the mass of the calorimeter, c_{cal} its specific heat capacity (typically copper, $c_{cal} = 400\,\text{J kg}^{-1}\text{K}^{-1}$).

Writing the heat balance:

$$m_1 c_1 (T_1 - T_f) = m_2 c_2 (T_f - T_2) + m_{cal} c_{cal} (T_f - T_2)$$

The left-hand side is the heat lost by the hot substance as it cools from T_1 to T_f. The right-hand side is the heat gained by the water plus the heat gained by the calorimeter, both of which warm from T_2 to T_f. Solving for c_1:

$$\mathbf{c_1 = \frac{(m_2 c_2 + m_{cal} c_{cal})(T_f - T_2)}{m_1 (T_1 - T_f)}}$$

The experimental arrangement is shown in the following figure. The calorimeter sits inside an insulating jacket to minimise heat loss to the surroundings. A thermometer measures the water temperature continuously. A stirrer keeps the water well mixed so that the thermometer reads a uniform temperature rather than a hot spot near the dropped block.

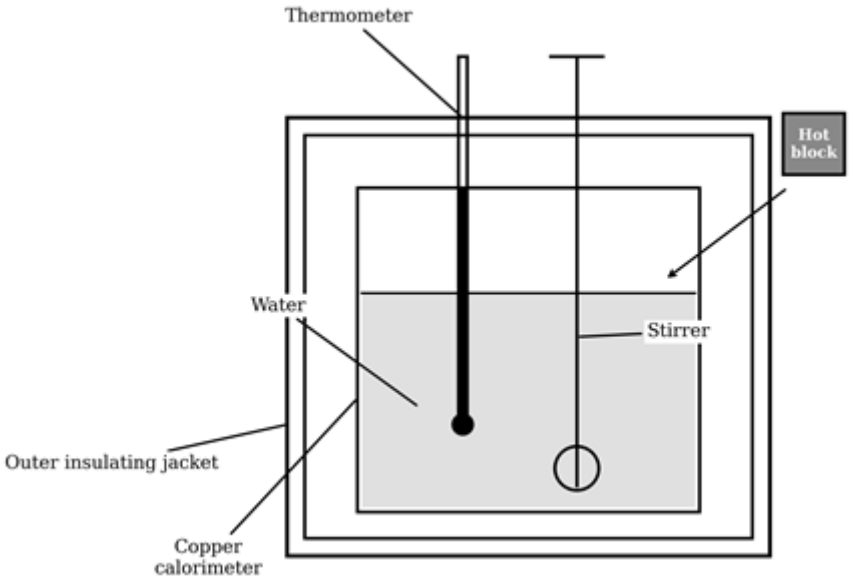


Figure: ***The method of mixtures.*** *A calorimeter containing water sits inside an outer insulating jacket. A thermometer records the water temperature, a stirrer keeps the water well mixed, and the hot substance (shown here as a metal block) is lowered into the water to allow heat to flow from it to the surrounding water and calorimeter until equilibrium is reached.*

Two sources of systematic error limit the accuracy of the method of mixtures.

1) **Heat loss from the calorimeter to the surroundings** during the experiment, which removes some of the heat that should have gone into warming the water. This is minimised by using an insulating jacket and by working quickly.
2) **Incomplete mixing,** in which the water near the hot block reaches a higher local temperature than the water at the thermometer and the thermometer reads below the true equilibrium value. This is minimised by thorough stirring.

A third, smaller source of error is the heat absorbed by the thermometer and stirrer themselves, which is usually neglected for rough work and corrected by an additional water-equivalent term for precision work.

Measuring specific heat capacity: the electrical method

The method of mixtures works well for solids but is awkward for liquids, because heating a liquid to a known high temperature in one vessel and transferring it without loss into another is very difficult. For liquids, the **electrical heating method** is usually preferred. A heater of known electrical power is immersed in the liquid inside an insulated calorimeter, switched on for a measured time, and the temperature rise is recorded. The heat supplied by the electrical heater equals the electrical energy dissipated:

$$Q = Pt = VIt$$

where V is the voltage across the heater, I the current through it, and t the time for which it is switched on. Setting this equal to the heat absorbed by the liquid and calorimeter:

$$VIt = \left(m_{liq}c_{liq} + m_{cal}c_{cal}\right)\Delta T$$

Solving for c_{liq}:

$$\mathbf{c_{liq} = \frac{VIt - m_{cal}c_{cal}\,\Delta T}{m_{liq}\,\Delta T}}$$

The experimental arrangement is shown in the following figure. An ammeter in series with the heater measures the current, a voltmeter across the heater measures the voltage, and the product VI is the power dissipated in the heater.

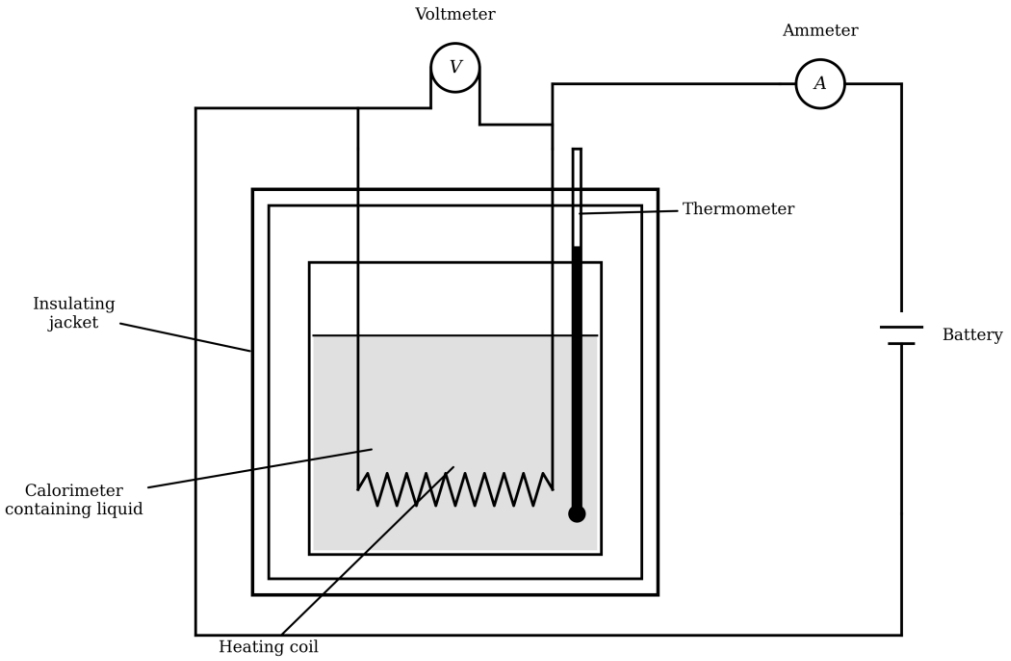


Figure: *The electrical heating method for measuring the specific heat capacity of a liquid. The heating coil is immersed in the liquid and is powered by a battery through an ammeter (in series) and a voltmeter (in parallel with the coil). The thermometer records the rise in temperature of the liquid, and the total electrical energy supplied by the coil is calculated as VIt.*

The electrical method is more accurate than the method of mixtures for two reasons. First, the electrical energy supplied can be measured with much greater precision than the heat transferred during a physical mixing. Second, the heater is inside the liquid throughout the experiment, so no transfer of a hot body is required and there is no risk of losing heat during the transfer.

Applications of specific heat capacity

Four familiar phenomena all trace back to specific heat capacity, each one more revealing than the last.

1. Cooking utensils

A cook in a hurry wants a pot that heats up quickly, and she therefore prefers a utensil with a low specific heat capacity (and low mass, since what matters is the thermal capacity $C = mc$). An aluminium sufuria ($c = 900\text{J kg}^{-1}\text{K}^{-1}$) warms up much faster than an equally heavy cast-iron sufuria ($c = 450\text{J kg}^{-1}\text{K}^{-1}$) when both are placed over the same flame, because the aluminium absorbs less heat per degree even though its specific heat capacity is higher; what matters is that the aluminium *sufuria* is typically much thinner and lighter, so its thermal capacity is lower overall. The flip side of the same principle is that a cook who wants food to stay hot during a long meal prefers a high thermal capacity, which keeps the pot warm for longer after the flame is switched off. A heavy cast-iron pot retains heat for an hour after cooking ends; a thin aluminium pan cools in minutes.

2. Water as a coolant

Water is used as the coolant in car engines, nuclear reactors, and power stations because its specific heat capacity of $4200\text{J kg}^{-1}\text{K}^{-1}$ is the highest of any common liquid. A kilogram of water flowing through an engine can absorb large amounts of heat for only a small rise in temperature, so a relatively small mass of water at a relatively low flow rate is enough to keep the engine from overheating. No other fluid in the cook's or engineer's everyday supply comes close. Oil, used where water cannot be used because of its freezing point or its electrical conductivity, has a specific heat capacity of about $2000\text{J kg}^{-1}\text{K}^{-1}$, which is less than half that of water.

3. Fish in frozen ponds

On a cold night in the highlands, the surface of a pond can freeze over while the fish beneath survive the winter. Three properties of water collaborate to make this possible. **The first** is water's large specific heat capacity, which means the pond takes a very long time to cool, and any heat stored in the lower layers is released slowly as the surface cools. **The second** is water's anomalous expansion near freezing: water reaches its maximum density at 4°C, not at 0°C, so the densest water settles to the bottom at 4°C and colder water floats above it. **The third** is that ice, once it forms on the surface, has a much lower thermal conductivity than water (we will meet thermal conductivity in Chapter 3), so the ice layer acts as an insulating blanket between the cold air above and the 4°C water beneath. The fish live in the 4°C reservoir at the bottom of the pond, insulated from the freezing air by the ice above and from the falling temperature by water's enormous capacity to store heat.

4. Sea breeze and land breeze

A coastal town experiences a distinctive daily pattern: a cool breeze blows in from the sea during the day, and a warmer breeze blows out toward the sea at night. The cause is the difference in effective thermal capacity between sea and land. The sea absorbs the sun's heat into a layer many metres deep, mixed by currents and waves, and warms very slowly because water's specific heat capacity is large and a large mass is involved. The land absorbs heat only into its top few centimetres and warms much more quickly. By midday, the land is hotter than the sea, and the air above the hot land rises; cooler air flows in from the sea to replace it, producing the sea breeze. At night, the land cools more quickly than the sea, the situation reverses, and a land breeze blows out toward the sea. Coastal towns are therefore milder year-round than inland towns at the same latitude, because the nearby sea acts as a giant thermal reservoir that warms slowly in summer and cools slowly in winter.

Theory and you have just defined specific heat capacity, measured it two different ways, and watched it power four familiar phenomena. Time for thermal equilibrium with three worked examples to put the formulas to work.

BINDER Example 15

An electric kettle rated 2.5kW contains 1.0L of water at 25°C. Taking the specific heat capacity of water as $4200\text{J kg}^{-1}\text{K}^{-1}$ and the density of water as 1000kg m^{-3}, calculate the time the kettle needs to heat the water to 100°C, assuming no heat is lost to the surroundings.

Solution

The mass of water is $m = \rho V = 1000\text{kg m}^{-3} \times 1.0 \times 10^{-3}\text{m}^3 = 1.0\text{kg}$.

The heat required to raise the temperature of the water from 25°C to 100°C is

$$\Delta Q = mc\Delta T = 1.0\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times 75\text{K} = 3.15 \times 10^5\text{J}$$

The kettle supplies electrical energy at a rate of $P = 2.5\text{kW} = 2500\text{J s}^{-1}$. The time required is therefore

$$t = \frac{\Delta Q}{P} = \frac{3.15 \times 10^5\text{J}}{2500\text{J s}^{-1}} = 126\text{s} \approx 2.1\text{minutes}$$

Making Sense of the Answer: *Two minutes is about how long a* 2.5kW *kettle actually takes to boil a litre of water in practice, which is a reassuring confirmation that the calculation is reasonable. Real kettles lose some heat to the surroundings and have their own thermal capacity to warm up, so the actual time is typically a few percent longer than our calculation suggests, but not dramatically so. The dominant physics is the heating of the water itself, and the corrections are secondary.*

Think Like a Physicist: *Any time a problem specifies a power and asks for a time, or specifies a time and asks for a power, the quantity linking them is the total heat* $\Delta Q = mc\Delta T$. *The heat is always computed first from the thermal properties of the substance being heated, and the time or power is then extracted by simple division. The calculation is exactly the same for a*

kettle heating water, a heater warming a room, or a rocket engine warming its exhaust: three quantities, one equation, one division.

BINDER Example 16

A copper calorimeter of mass 100g contains 200g of water at 15°C. A 150g block of iron, initially at 100°C, is dropped into the water. Assuming no heat is lost to the surroundings and taking $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$, $c_{Fe} = 450\text{J kg}^{-1}\text{K}^{-1}$, and $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$, calculate the final equilibrium temperature.

Solution

Let T_f be the final equilibrium temperature in °C. Applying the heat balance (heat lost by iron equals heat gained by water plus heat gained by calorimeter):

$$m_{Fe}c_{Fe}(T_1 - T_f) = m_{water}c_{water}(T_f - T_2) + m_{Cu}c_{Cu}(T_f - T_2)$$

Substituting the given values in SI units:

$$(0.150)(450)(100 - T_f) = (0.200)(4200)(T_f - 15) + (0.100)(400)(T_f - 15)$$

$$67.5(100 - T_f) = 840(T_f - 15) + 40(T_f - 15)$$

$$67.5(100 - T_f) = 880(T_f - 15)$$

$$T_f = \frac{19950}{947.5} = 21.1°\text{C}$$

The final equilibrium temperature is approximately 21.1°C.

Making Sense of the Answer: *The iron block cools by nearly* $80°C$ *(from* 100 *down to* 21.1*) while the water warms by only* $6°C$ *(from* 15 *up to* 21.1*). The water + calorimeter has about thirteen times the thermal capacity of the iron, so the iron's temperature change is about thirteen times the water's.*

Think Like a Physicist: *A method-of-mixtures problem is always a heat-balance equation: the heat lost by the hot body equals the heat gained by everything else. Write the equation carefully before substituting numbers, listing every object that changes temperature (the hot substance, the liquid, the calorimeter, and any other solid in thermal contact). If the calorimeter is made of the same material as the hot substance, you can still include it; its term simply has a different mass and a different initial temperature. Never drop the calorimeter term without checking whether it matters; for small hot masses and large calorimeters, the calorimeter's contribution can dominate the right-hand side.*

REAL Example 17

Kipanga has just come home from the market carrying two things: a new aluminium sufuria that his mother wanted for cooking ugali, and a sharp pedagogical disagreement with their neighbour Msemakweli. Mr. Akilikubwa is repairing his gate next door, and Kipanga stops to complain.

Kipanga: *"Sir, my mother sent me to buy an aluminium sufuria, saying it heats up fast so she can cook ugali quickly. But Msemakweli saw me walking with it and laughed. He said his cast-iron sufuria is much better because it keeps the food warm for hours after he switches off the stove. Now my mother and Msemakweli cannot both be right, can they? One of them has to be wrong about which sufuria is superior."*

Mr. Akilikubwa sets down his hammer, looks at the shiny aluminium sufuria Kipanga is holding up, and smiles.

He then clearly explained the following to Kipanga: (a) why Kipanga's mother is correct that the aluminium sufuria heats up quickly, (b) why Msemakweli is correct that the cast-iron sufuria stays warm longer, and (c) why both can be right at the same time without contradicting each other.

What was Mr. Akilikubwa's explanation?

Solution

(a) Kipanga's mother is correct because the aluminium sufuria has a low thermal capacity. An aluminium sufuria of the typical thickness sold in a Tanzanian market is both thin-walled and light, which means its mass is small. Although aluminium's specific heat capacity ($900\text{J kg}^{-1}\text{K}^{-1}$) is actually twice that of iron ($450\text{J kg}^{-1}\text{K}^{-1}$), the mass of the aluminium pot is so much smaller than the mass of an equivalent cast-iron pot that the product $C = mc$, the thermal capacity of the whole pot, is lower for the aluminium. A smaller thermal capacity means less heat is needed to warm the pot itself before the pot starts warming the food inside it, so the burner's heat goes into the food quickly and the ugali cooks fast.

(b) Msemakweli is correct because the cast-iron sufuria has a large thermal capacity. A cast-iron pot is heavy (high density) and thick-walled, so its mass is large. Even though iron's specific heat capacity is lower than aluminium's, the huge mass more than makes up for it, and the thermal capacity $C = mc$ of the cast-iron pot is much larger than that of the aluminium pot. A pot with a large thermal capacity holds a large amount of heat at a given temperature, so when the heat source is off the pot continues to release this stored heat into the food for a long time before its own temperature drops significantly. The ugali therefore stays hot longer after cooking, not because the pot is a better conductor but because it stores more heat.

(c) Both observations are correct because they are about two different physical quantities that happen to point in opposite directions for the two materials. Kipanga's mother is commenting on how fast the pot reaches cooking temperature, which is governed by the thermal capacity (a small C means fast heating). Msemakweli is commenting on how long the pot keeps its contents warm, which is also governed by the thermal capacity (a large C means slow cooling). The two statements describe opposite virtues of different pots, and both are right within their own context. Neither speaker is wrong; they simply want different things from their pots, and the market has provided them with a range of choices that trade one virtue for the other.

The physics does not force us to pick a winner between the two sufurias. Instead, it tells us what each pot is good for. A small aluminium sufuria is ideal for cooking ugali quickly in the morning when time is short. A heavy cast-iron pot is ideal for slow stews and for keeping food warm during a long meal where many guests arrive at different times. A wise cook owns both and uses each for what it does best.

Making Sense of the Answer: *The apparent paradox in Kipanga's question dissolves as soon as we distinguish between two closely related ideas. Fast heating and slow cooling are not opposites; they are two consequences of the same underlying property, the thermal capacity. A pot with low* C *heats fast and cools fast; a pot with high* C *heats slowly and cools slowly. Which is better depends entirely on what the cook needs. The disagreement between Kipanga's mother and Msemakweli is therefore not a disagreement about physics but a disagreement about kitchen strategy, and both of them are correct within their own strategies.*

Think Like a Physicist: *Whenever you hear two people disagree about which of two objects is "better" for some thermal purpose, ask what each speaker really means by "better".*

Specific heat capacity has taken us a long way. Every calculation in the examples above shared a common assumption, namely that warming a substance by one degree requires a certain quantity of heat, and that this quantity of heat scales in a simple linear way with both the mass and the temperature rise. But the assumption hides something that any cook already knows. When water at $100°C$ is heated further, it does not simply continue to get hotter. Instead, it boils. During the boiling, vast amounts of heat continue to flow into the water, and yet the thermometer refuses to climb above $100°C$. *Where is all the heat going, if it is not raising the temperature?* The next section answers this question with a new and different quantity, the latent heat, which accounts for the energy absorbed or released whenever a substance changes from one phase of matter to another. Get the kettle back on the stove. We have one more act of thermal physics to watch unfold.

LATENT HEAT

Put a saucepan of ice cubes on a flame and watch a thermometer in the ice. At first the temperature climbs steadily as the ice warms, following exactly the specific-heat-capacity story of the previous section. At zero degrees Celsius, the climbing stops. The flame is still on. The saucepan is still getting heat. The thermometer, placed in the melting ice, reads stubbornly zero degrees for minute after minute. Only after the very last piece of ice has melted does the thermometer begin to rise again, and the water in the saucepan heats towards the boiling point. At one hundred degrees Celsius, exactly the same thing happens. The temperature stops climbing, the water boils steadily, the thermometer refuses to move above one hundred degrees, and the saucepan slowly empties into the steam above. Only when the last drop of water has boiled away does the steam, if contained, begin to rise in temperature again.

Where did the heat go during the two flat periods? The flame never stopped. The saucepan never cooled. Yet the thermometer sat still for minutes at a stretch, as if the flame had temporarily been switched off. The answer is that the heat did not disappear. It was used to change the substance from one phase of matter to another, from ice to liquid water at zero degrees and from liquid water to steam at one hundred degrees. A phase change requires a specific amount of energy per kilogram, and during the phase change that energy is absorbed without raising the temperature of the substance at all. The quantity of energy required is called the **latent heat**.

Defining latent heat

The latent heat of a phase change is defined as *the heat required, per unit mass, to change a substance from one phase to another at the phase-change temperature without changing the temperature itself.*

$$L = \frac{Q}{m} \qquad \Rightarrow \qquad Q = mL$$

The SI unit is J kg^{-1}. Two different latent heats appear in ordinary thermal physics:

- **Latent heat of fusion**, L_f, *the heat required to melt a kilogram of a solid at its melting point into a kilogram of liquid at the same temperature.*
- **Latent heat of vapourisation**, L_v, *the heat required to vapourise a kilogram of a liquid at its boiling point into a kilogram of vapour at the same temperature.*

For water, the two latent heats are:

$$L_f = 3.34 \times 10^5 \text{J kg}^{-1} \qquad L_v = 2.26 \times 10^6 \text{J kg}^{-1}$$

The vapourisation value is about seven times larger than the fusion value. A single kilogram of water at one hundred degrees requires over two million joules of heat to boil away completely, enough energy to lift a one-tonne lorry a vertical height of

about two hundred and thirty metres! The fusion of a single kilogram of ice at zero degrees requires considerably less, but still an amount comparable to the heat needed to raise the same kilogram of water by eighty degrees.

Why L_v is so much larger than L_f

The molecular explanation of the two latent heats is simple and physical, and it explains the factor of seven cleanly. In the solid phase, water molecules are held rigidly in place by strong intermolecular forces and can vibrate only about their fixed positions. In the liquid phase, the molecules are still close to each other but are free to slide past one another; the intermolecular forces have been loosened, not broken. In the vapour phase, the molecules are widely separated, the intermolecular forces have been completely broken, and the molecules move independently in all directions.

Melting takes a kilogram of ice and turns it into a kilogram of liquid water. The intermolecular forces are only partially broken, because the molecules of the liquid are still in contact with each other most of the time. The heat required to accomplish this partial loosening is L_f.

Boiling takes a kilogram of liquid water and turns it into a kilogram of vapour. Now the intermolecular forces must be broken completely, and the molecules must be given enough kinetic energy to fly apart and occupy a volume roughly 1700 times larger than the liquid they left behind. A further contribution comes from the external work done by the expanding vapour against the atmospheric pressure pushing down on it. All of this energy must be supplied by the heat source, and the total is L_v.

The factor of seven between L_v and L_f is therefore the energy cost of the difference between loosening the intermolecular forces and breaking them entirely. If water molecules were held together only weakly, L_f and L_v would be closer in size. If they were held together very strongly, both latent heats would be larger and so would their ratio. The observed ratio of about seven is a direct fingerprint of how tightly the water molecule clings to its neighbours.

The heating curve for water

All of the ideas above are summarised in a single graph: the **heating curve** for water, showing the temperature of one kilogram of ice as heat is supplied steadily, from a starting temperature of $-20°\text{C}$ until the water has been fully converted to steam at $120°\text{C}$.

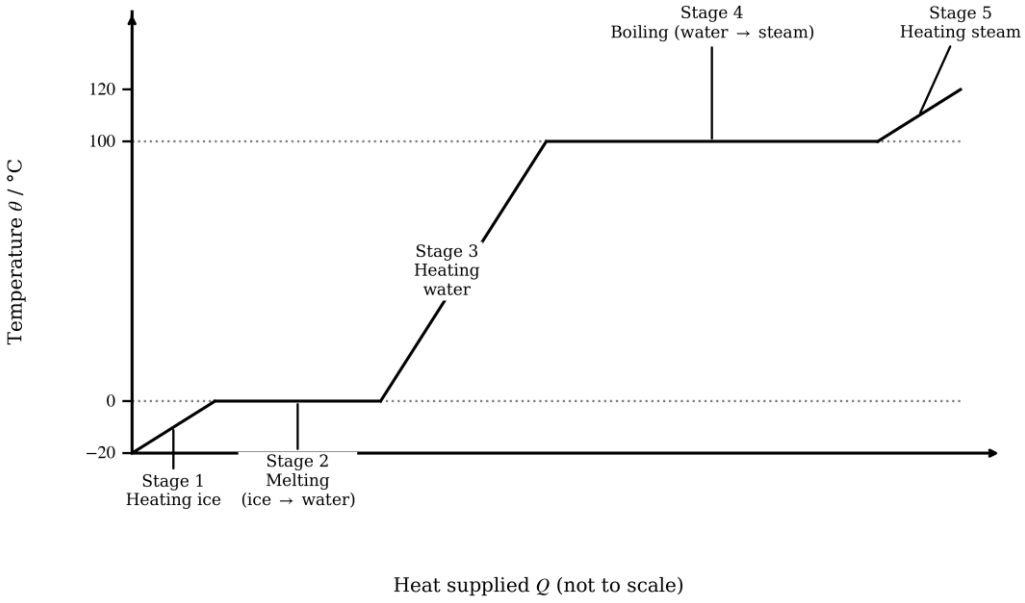


Figure: ***The heating curve for one kilogram of ice, warmed steadily from $-20°C$ to $120°C$.*** *Five stages are visible. Stage 1: the ice warms from $-20°C$ to $0°C$, a sloped segment whose slope depends on the specific heat capacity of ice. Stage 2: the ice melts at $0°C$, a flat segment whose width is L_f. Stage 3: the liquid water warms from $0°C$ to $100°C$, a sloped segment whose slope depends on the specific heat capacity of liquid water. Stage 4: the water boils at $100°C$, a flat segment whose width is L_v. Stage 5: the steam warms from $100°C$ to $120°C$, a sloped segment whose slope depends on the specific heat capacity of steam. The width of Stage 4 is much larger than the width of Stage 2, because L_v is about seven times L_f.*

The three sloped segments and the two flat segments each tell a different story. On a sloped segment, heat is being used to raise the temperature of a substance that is not changing phase. The slope $\mathrm{d}\theta/\mathrm{dQ}$ is inversely proportional to the specific heat capacity, $\frac{\mathrm{d}\theta}{\mathrm{dQ}} = \frac{1}{\mathrm{mc}}$, so steeper slopes correspond to substances with smaller c. The slope of Stage 5 (heating steam) is steeper

than the slope of Stage 3 (heating water) because $c_{steam} \approx 2000 \text{J kg}^{-1}\text{K}^{-1}$ is less than half of $c_{water} = 4200 \text{J kg}^{-1}\text{K}^{-1}$, so a given amount of heat raises the steam's temperature more than twice as much as it raises the water's.

On a flat segment, heat is being absorbed but the temperature is not rising. All of the incoming heat is going into breaking or loosening intermolecular bonds, a process that changes the molecular potential energy of the substance without changing the molecular kinetic energy (and temperature is, at root, a measure of average molecular kinetic energy). The width of the flat segment is equal to the total heat absorbed during the phase change, which is mL_f for melting and mL_v for boiling. The ratio of the two widths is the ratio $\frac{L_v}{L_f} \approx 6.8$, which is why Stage 4 in the figure is visibly much wider than Stage 2.

Evaporation and boiling are not the same thing

A common confusion is between evaporation and boiling. They both turn liquid water into water vapour, but they are very different processes, and the difference matters.

Evaporation *occurs at the free surface of a liquid, at any temperature above its freezing point.* The molecules at the surface have a distribution of kinetic energies, and a small fraction of them at any instant have enough energy to escape into the air above. Those escaping molecules carry away kinetic energy from the liquid, which cools the liquid slightly (this is the cooling effect of evaporation, familiar from sweating skin and wet laundry hanging in the sun). Evaporation is a slow process, limited to the surface, and driven by the random thermal motion of molecules at that surface. It does not require the liquid to reach the boiling temperature; a puddle on the ground at $25°\text{C}$ will evaporate away in a few hours.

Boiling *occurs throughout the volume of a liquid, at the specific temperature where the vapour pressure of the liquid equals the atmospheric pressure above it.* At this temperature, bubbles of vapour can form anywhere inside the liquid and rise to the surface; the liquid visibly churns. Boiling is a fast process, not limited to the surface, and it requires heat to be supplied from outside the liquid to keep the phase change going. Boiling temperature depends on atmospheric pressure: at sea level, water boils at $100°\text{C}$; on a mountain top where the atmospheric pressure is lower, water boils at a lower temperature; and in a pressure cooker where the pressure is higher, water boils at a temperature above $100°\text{C}$.

Both processes absorb the same latent heat of vapourisation, because both convert liquid to vapour and the molecular cost per molecule is the same. The difference is entirely in where and how fast the phase change happens.

The first law of thermodynamics and phase changes

When a kilogram of water at $100°\text{C}$ is boiled away into steam at $100°\text{C}$ and standard atmospheric pressure, the heat absorbed is $Q = L_v = 2.26 \times 10^6 \text{J}$ per kilogram. Not all of this heat ends up inside the steam as internal energy. Part of it is used by the steam to push back the atmosphere as it expands from the volume of the water (about one litre per kilogram) to the volume of the steam (about 1670 litres per kilogram). The work done by the steam against the atmospheric pressure is $W = P\Delta V$, where ΔV is the increase in volume.

The remaining energy stays inside the steam as an increase in its internal energy, ΔU. The three quantities are related by the first law of thermodynamics,

$$\boldsymbol{\Delta U = Q - W}$$

which we will derive carefully in Chapter 5. For the present chapter, it is enough to know that the first law says the heat absorbed is partitioned into two contributions: some raises the internal energy of the substance, and the rest does external work against the surroundings. For a phase change from liquid to vapour at constant atmospheric pressure, W is the work done by the expanding vapour and ΔU is the part of Q that actually ends up as internal molecular energy of the steam.

Heat conduction and the ice pond

One final application of latent heat is the growth of an ice layer on the surface of a frozen pond. When the air above a pond drops below $0°\text{C}$, the water at the surface freezes to form an ice layer, and the ice layer grows thicker over time as more water freezes at its lower surface. The rate at which the layer grows is controlled by two physical processes: the rate at which heat is conducted through the ice layer from the water below to the cold air above, and the rate at which the water at the bottom of the ice layer freezes.

The rate of heat conduction through a layer of thickness **x** with cross-section area **A**, thermal conductivity **k**, and a temperature difference **ΔT** across it is given by Fourier's law:

$$\boldsymbol{\frac{dQ}{dt} = \frac{kA\Delta T}{x}}$$

We will derive Fourier's law carefully in Chapter 3, but for now we need it as a working tool. The heat conducted through the ice is the heat released by the water beneath as it freezes, and if an additional thin layer of water of thickness dx freezes, the heat released is $dQ = \rho A\, dx\, L_f$, where ρ is the density of water (approximately equal to the density of ice, for this calculation).

Setting the rate of heat conducted equal to the rate at which this freezing-heat is produced:

$$\frac{kA\Delta T}{x} = \rho A\, L_f \frac{dx}{dt}$$

The area A cancels from both sides:

$$\frac{k\Delta T}{x} = \rho L_f \frac{dx}{dt} \quad \Rightarrow \quad \mathbf{\frac{dx}{dt} = \frac{k\Delta T}{\rho L_f x}}$$

This equation tells us that the growth rate of the ice layer is inversely proportional to the layer's current thickness. When the layer is thin, the ice grows quickly; as the layer thickens, the growth slows because the same temperature difference is driving heat through a thicker insulating barrier. This is why a thin ice layer on a pond is untrustworthy (it is still growing fast and its thickness changes rapidly), while a thick layer is stable (its growth has slowed to a crawl). Worked examples later in this section will put concrete numbers into this equation.

Theory and you have just defined latent heat, explained why vapourisation requires so much more energy than fusion, traced a full heating curve, and met two important forward-references (the first law for phase changes and Fourier's law for heat conduction). Time for thermal equilibrium with five worked examples.

BINDER Example 18

A 200g sample of an unknown solid is heated steadily from $-40°C$ through its melting point at $40°C$ up to a final temperature of $80°C$. A plot of temperature against total heat supplied shows three distinct segments. The first segment (heating the solid) absorbs $3600J$ and spans the temperature range $-40°C$ to $40°C$. The second segment is a horizontal region at $40°C$, absorbing $16000J$ without any temperature change. The third segment (heating the liquid) absorbs $4000J$ and spans $40°C$ to $80°C$. Calculate (a) the specific heat capacity of the solid phase, (b) the specific latent heat of fusion, (c) the specific heat capacity of the liquid phase.

Solution

Using $m = 0.200kg$ throughout.

(a) The first segment is a sloped region in which $3600J$ raises the solid's temperature by:

$$\Delta T_1 = 40℃ - (-40℃) = 80°C$$

Applying $Q = mc\Delta T$ and solving for c:

$$c_{solid} = \frac{Q_1}{m\,\Delta T_1} = \frac{3600J}{0.200kg \times 80℃} = 225J\ kg^{-1}℃^{-1}$$

The specific heat capacity of the solid phase is $225J\ kg^{-1}℃^{-1}$.

(b) The second segment is a horizontal region at the melting point, in which $16\,000J$ is absorbed at constant temperature. Applying $Q = mL_f$ and solving for L_f:

$$L_f = \frac{Q_2}{m} = \frac{16000J}{0.200kg} = 8.0 \times 10^4 J\ kg^{-1}$$

The specific latent heat of fusion is $8.0 \times 10^4 J\ kg^{-1}$.

(c) The third segment is a sloped region in which $4000J$ raises the liquid's temperature by:

$$\Delta T_3 = 80℃ - 40℃ = 40°C$$

Applying $Q = mc\Delta T$ again:

$$c_{liquid} = \frac{Q_3}{m\,\Delta T_3} = \frac{4000J}{0.200kg \times 40℃} = 500J\ kg^{-1}℃^{-1}$$

The specific heat capacity of the liquid is $500J\ kg^{-1}℃^{-1}$.

Making Sense of the Answer: *The three numerical values are all smaller than the corresponding values for water. The solid's specific heat capacity is about one-tenth of ice's, the latent heat of fusion is about one-quarter of ice's, and the liquid's specific heat capacity is about one-eighth of water's. The substance is therefore nothing like water; it might be a metal such as lead or bismuth, for which the quoted values are roughly the right order of magnitude. A heating curve, read carefully, tells the physicist not only how a substance behaves thermally but also gives clues about what kind of substance it is.*

Think Like a Physicist: *A heating curve can be read in either direction. Given a substance and its thermal properties, you can sketch the curve; given a curve and its numerical segment sizes, you can recover the thermal properties. The two directions are the same physics seen from opposite sides. Whenever a problem gives you energies absorbed on specific segments of a heating curve, identify which segment corresponds to which phase or phase change, and apply either* $Q = mc\Delta T$ *(for sloped segments) or* $Q = mL$ *(for flat segments) with the appropriate value of* ΔT *or* m.

BINDER Example 19

Calculate the total heat required to convert $500g$ of ice at $-10°C$ into steam at $100°C$. Take $c_{ice} = 2100J\ kg^{-1}K^{-1}$, $c_{water} = 4200J\ kg^{-1}K^{-1}$, $L_f = 3.34 \times 10^5 J\ kg^{-1}$, $L_v = 2.26 \times 10^6 J\ kg^{-1}$.

Solution

The process has four distinct stages, each contributing a separate amount of heat. Let $m = 0.500kg$ throughout.

Stage 1. *Heating the ice from* $-10°C$ *to* $0°C$*:*

$$Q_1 = m\,c_{ice}\,\Delta T = 0.500kg \times 2100J\ kg^{-1}K^{-1} \times 10K = 1.05 \times 10^4 J$$

Stage 2. *Melting the ice at* $0°C$ *into water at* $0°C$*:*

$$Q_2 = m\,L_f = 0.500\text{kg} \times 3.34 \times 10^5\text{J kg}^{-1} = 1.67 \times 10^5\text{J}$$

Stage 3. *Heating the water from* $0°C$ *to* $100°C$*:*

$$Q_3 = m\,c_{water}\,\Delta T = 0.500\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times 100\text{K} = 2.10 \times 10^5\text{J}$$

Stage 4. *Boiling the water at* $100°C$ *into steam at* $100°C$*:*

$$Q_4 = m\,L_v = 0.500\text{kg} \times 2.26 \times 10^6\text{J kg}^{-1} = 1.13 \times 10^6\text{J}$$

The total heat is the sum of all four contributions:

$$Q_{total} = Q_1 + Q_2 + Q_3 + Q_4$$

$$Q_{total} = (1.05 \times 10^4 + 1.67 \times 10^5 + 2.10 \times 10^5 + 1.13 \times 10^6)\text{J} = 1.52 \times 10^6\text{J} = 1.52\text{MJ}$$

Making Sense of the Answer: *Three-quarters of the total energy went into Stage 4, the vapourisation of the water, even though that stage did not change the temperature by a single degree. This is a direct consequence of* L_v *being so much larger than every other contribution in the calculation. If you were designing a solar cooker or a gas boiler, knowing that roughly seventy-five per cent of the energy needed to go from ice to steam is spent during boiling alone is essential, because it tells you where your appliance's heat budget actually goes.*

Think Like a Physicist: *Whenever a multi-stage heat calculation crosses one or more phase changes, break the calculation into separate stages before attempting the arithmetic. One stage for each segment of the heating curve. The specific heat capacity applies on sloped segments, the latent heat applies on flat segments, and the total heat is the sum of the contributions from all stages. Working this way makes it almost impossible to forget a phase change or use the wrong specific heat.*

REAL Example 20

In rural areas of **Tanzania** where electric fans and air conditioners are not available, a traditional remedy for a feverish child is to wipe the child's forehead with a damp cloth. The cloth has been dipped in cool water and wrung out, and it is left resting on the child's forehead until it dries out, at which point it is dipped again. Parents observe that this simple procedure brings the fever down noticeably within minutes, while simply fanning the child with a dry hand has much less effect. Using the concepts of specific heat capacity and latent heat of vapourisation, explain why the damp cloth is so effective at cooling the child, and estimate the quantity of heat removed from the forehead when 5g of water evaporates from the cloth. Take $L_v = 2.26 \times 10^6\text{J kg}^{-1}$ for water at body temperature and $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$.

Solution

The cloth cools the forehead by two separate mechanisms, one of which is much more powerful than the other.

The first mechanism is straightforward *temperature equilibration*. The cool water in the cloth is initially at, say, $25°\text{C}$, while the forehead is at about $38°\text{C}$ in a feverish child. Heat flows from the forehead to the water by conduction, warming the water and cooling the forehead. The heat that must be absorbed to raise 5g of water from $25°\text{C}$ to $38°\text{C}$ is

$$Q_1 = mc\,\Delta T = 0.005\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times 13\text{K} = 273\text{J}$$

This contribution, about 273 joules, accounts for the cloth feeling cool to the touch when first applied, but it is not the dominant cooling mechanism.

The second mechanism is *evaporation.* Once the water in the cloth has warmed to body temperature, it begins to evaporate from the surface of the cloth into the air. Each molecule of water that escapes into the air carries away energy equal to the latent heat of vapourisation per molecule, and that energy is drawn from the forehead. The heat absorbed when all 5g of water in the cloth evaporates is

$$Q_2 = mL_v = 0.005\text{kg} \times 2.26 \times 10^6\text{J kg}^{-1} = 1.13 \times 10^4\text{J}$$

This is about 11300J, more than forty times larger than the temperature-equilibration contribution. The damp cloth is an effective cooler because evaporation of water is a hugely energy-expensive process, and every gram of water that evaporates from the cloth draws over two thousand joules of heat out of whatever surface it was in contact with.

Fanning with a dry hand, by contrast, only accelerates the loss of heat by convection. It does not trigger any evaporation from the skin unless the skin is already sweating, and even then the cooling is limited to the small amount of sweat that is actually on the skin. The damp cloth brings a fresh supply of water to the forehead and holds it in intimate contact with the skin long enough for nearly all of it to evaporate, delivering cooling equivalent to forty or fifty times what a dry fan could provide.

The total heat removed from the forehead by the damp cloth is therefore

$$Q_{total} = Q_1 + Q_2 = 273\text{J} + 1.13 \times 10^4\text{J} \approx 1.16 \times 10^4\text{J}$$

Making Sense of the Answer: *The dominant contribution is the evaporation term, and this is not an accident. Every cooling system that relies on water, from the sweat on a marathon runner's skin to the cooling towers of a power station, exploits the same huge latent heat of vapourisation. The reason parents have trusted the damp-cloth remedy for generations is that the physics is on their side: water's latent heat of vapourisation is one of the largest in nature, and evaporative cooling is a far more effective heat-removal mechanism than ordinary conduction or convection at the same temperatures.*

Think Like a Physicist: *Whenever a cooling process involves water and the temperatures are near body temperature, check whether evaporation is occurring. If it is, the latent heat contribution almost always dominates over the specific heat contribution, because* L_v *for water is numerically much larger than* c_{water} *multiplied by any reasonable temperature change. This is why sweating is effective at cooling the body, why wet washing dries faster on a windy day, and why spilling water on a hot stove cools the stove surface dramatically.*

HOT Example 21

A mass of 1.0g of water at $100°C$ is converted into steam at $100°C$ at the standard atmospheric pressure $P = 1.013 \times 10^5 Pa$. The volume of 1.0g of water is $1.0 cm^3$, and the volume of 1.0g of steam at $100°C$ and standard pressure is $1671 cm^3$. Take $L_v = 2.260 \times 10^6 J\ kg^{-1}$. Calculate:

(a) the heat absorbed during the phase change,

(b) the external work done by the expanding steam against the atmospheric pressure, and

(c) the increase in internal energy of the water.

Solution

(a) The heat absorbed during the phase change is $Q = mL_v$:

$$Q = 1.0 \times 10^{-3} kg \times 2.260 \times 10^6 J\ kg^{-1} = 2260 J$$

(b) The external work done by the steam against the atmospheric pressure is $W = P\Delta V$, where ΔV is the change in volume from liquid water to steam:

$$\Delta V = V_{steam} - V_{water} = 1671 cm^3 - 1.0 cm^3 = 1670 cm^3 = 1.670 \times 10^{-3} m^3$$

$$W = P\,\Delta V = 1.013 \times 10^5 Pa \times 1.670 \times 10^{-3} m^3 = 169.2 J \approx 169 J$$

(c) From the first law of thermodynamics, $\Delta U = Q - W$:

$$\Delta U = 2260 J - 169 J = 2091 J$$

Of the $2260 J$ absorbed, $169 J$ is spent on pushing back the atmosphere, and the remaining $2091 J$ goes into the internal energy of the steam.

Making Sense of the Answer: *The partition of the absorbed heat is about* 93% *to internal energy and* 7% *to external work. The internal-energy contribution is the dominant one, and it represents the energy stored in the widely separated steam molecules, which have broken free from their neighbours in the liquid. The external-work contribution is smaller but not negligible, and it is the reason the latent heat of vapourisation of water at atmospheric pressure is noticeably larger than it would be if the steam were allowed to form in a vacuum (where* W *would be zero and* ΔU *would equal* Q *directly).*

Think Like a Physicist: *Whenever a phase change occurs at constant atmospheric pressure, the heat absorbed is partitioned into two parts: one part pushes back the atmosphere (*$W = P\Delta V$*) and one part increases the internal energy (*$\Delta U = Q - W$*).* ***For condensed phases (solid to liquid),*** ΔV ***is small and*** W ***is negligible, so*** $\Delta U \approx Q$**.** *For phase changes involving a gas (liquid to vapour),* ΔV *is very large and* W *is appreciable; it must always be computed separately and subtracted from* Q *to get* ΔU*.*

HOT Example 22

On a very cold night in the Tanzanian highlands, a shallow pond develops a layer of ice on its surface. The air temperature above the ice is $-10°C$ and the water immediately beneath the ice is at $0°C$. Take the thermal conductivity of ice as $k = 2.2 W\ m^{-1}K^{-1}$, the density of ice as $\rho = 920 kg\ m^{-3}$, and the latent heat of fusion as $L_f = 3.34 \times 10^5 J\ kg^{-1}$. At the moment when the ice layer is 10cm thick, calculate:

(a) the rate of heat loss per square metre of the ice surface,

(b) the rate at which the thickness of the ice layer is increasing, and

(c) the time it would take for the ice layer to grow by an additional 1mm at this rate.

Solution

(a) The rate of heat flow per unit area through the ice, using Fourier's law $\frac{dQ}{dt} = \frac{kA\Delta T}{x}$ and dividing both sides by A:

$$\frac{1}{A}\frac{dQ}{dt} = \frac{k\,\Delta T}{x} = \frac{2.2 W\ m^{-1}K^{-1} \times 10K}{0.10m} = 220 W\ m^{-2}$$

The rate of heat loss is $220 Wm^{-2}$ at this instant.

(b) The growth rate of the ice layer comes from the relation derived earlier in this chapter,

$$\frac{dx}{dt} = \frac{k\,\Delta T}{\rho L_f x}$$

Substituting the numerical values:

$$\frac{dx}{dt} = \frac{2.2 W\ m^{-1}K^{-1} \times 10K}{920 kg\ m^{-3} \times 3.34 \times 10^5 J\ kg^{-1} \times 0.10m}$$

$$\frac{dx}{dt} = 7.16 \times 10^{-7}\,\text{m s}^{-1}$$

The thickness of the ice layer is increasing at the rate of $7.16 \times 10^{-7}\,\text{m s}^{-1}$.

This is about $0.72\,\mu\text{m s}^{-1}$, or equivalently $2.6\,\text{mm h}^{-1}$, or $62\,\text{mm d}^{-1}$ if the conditions remained constant.

(c) At the instantaneous rate computed in part (b), an additional 1mm of ice ($dx = 1 \times 10^{-3}\,\text{m}$) would form in a time

$$dt = \frac{dx}{\frac{dx}{dt}} = \frac{1 \times 10^{-3}\,\text{m}}{7.16 \times 10^{-7}\,\text{m s}^{-1}} = 1397\,\text{s} \approx 23\,\text{min}$$

It would take 1397s or approximately 23min.

This is slightly under twenty-four minutes, assuming the ice thickness does not change appreciably during this interval (which is a reasonable assumption since 1mm is only 1% of the current 10cm layer).

Making Sense of the Answer: *The ice grows at a rate of about 2.6mm per hour at the moment when the layer is 10cm thick. Once the layer reaches 20cm, the growth rate would drop to half that value, because $dx/dt \propto 1/x$. Once the layer reaches 1m thick (a rare but physically achievable state in Arctic conditions), the growth rate would drop by a factor of ten, explaining why even the thickest Arctic ice sheets take years rather than months to form. The inverse relationship between thickness and growth rate is the physics behind why a thick ice layer is stable and safe, while a thin ice layer is rapidly changing and treacherous.*

Think Like a Physicist: *Whenever a heat-conduction problem involves a moving phase boundary (ice growing into water, steam condensing onto a cold surface, frost forming on a window), the rate of phase change is governed by the balance between two physical processes: the rate at which heat is conducted away from the phase boundary, and the amount of heat per unit mass released (or absorbed) by the phase change itself. Set the two rates equal to each other and solve for the velocity of the boundary. The same technique works for every moving-boundary heat problem you will ever meet.*

With specific heat capacity and latent heat now behind us, every thermal quantity in Chapter 2 has been defined, measured, and tested against real physical situations. The calculations in the preceding sections covered kettles, calorimeters, cooking pots, clothes drying in the sun, ponds freezing over in the highlands, and the conversion of water into steam under an atmosphere of pressure. What remains is to step back and collect the various applications of thermal physics into a single narrative: where do these ideas reappear in the world of everyday life and engineering, and what other phenomena have we not yet touched that belong in the story? The next section is our closing tour of that terrain.

APPLICATIONS OF THERMAL PHYSICS

Every idea in this chapter has a job waiting for it in the world outside the classroom. Thermometers measure temperature in hospitals, factories, and kitchens. Specific heat capacity determines the size of a cooking pot, the flow rate of coolant in an engine, and the capacity of a solar water heater. Latent heat drives refrigerators, pressure cookers, steam autoclaves, and the sweating body. This closing section takes six of the most important applications of thermal physics and examines each in turn, showing how the quantities we have built fit together to solve real engineering and biological problems.

1. The refrigerator and the air conditioner

A refrigerator is a box that keeps its interior cooler than the room around it. Since heat naturally flows from hot to cold, the refrigerator has to actively pump heat the other way, from the cold interior to the warm exterior. The physics of this pumping is powered by latent heat of vapourisation.

Inside the cooling coils that run along the back or the underside of the refrigerator interior, a special working fluid is held at a pressure at which it boils at a low temperature, typically around $-25°\text{C}$. The fluid evaporates in these coils, drawing heat from the surrounding air inside the refrigerator as it does so, because evaporation requires latent heat L_v that must be supplied from somewhere. The air inside the refrigerator provides that heat, and in doing so the air cools, which is what keeps the food cold.

The vapour is then pumped by an electric compressor to the coils on the back of the refrigerator, where it is compressed to a higher pressure at which its boiling point rises above room temperature, say to 45°C. At this higher pressure, the vapour condenses back into liquid, releasing the same latent heat L_v it had absorbed inside. The heat released flows into the warmer room air, which is now cooler than the compressed vapour. The liquid then returns through an expansion valve to the interior coils, where the pressure drops, the boiling point falls back to $-25°\text{C}$, and the cycle repeats.

The refrigerator is therefore a clever use of the fact that a fluid's boiling point depends on pressure. By alternating the same fluid between two different pressures, the engineer arranges for the fluid to absorb heat at the low-pressure boiling point (cold side) and release it at the high-pressure boiling point (hot side). An air conditioner works identically; the only difference is that the "interior" is a room of the building and the "exterior" is the open air outside.

2. The pressure cooker

The reverse application of the same pressure-boiling-point relation is the pressure cooker. A sealed pot with a safety valve is heated on a flame. As the water inside boils, the steam produced cannot escape, and it accumulates above the water until the pressure inside the pot rises well above atmospheric. The elevated pressure raises the boiling point of the water, so the water continues to boil at a temperature higher than $100°C$.

In a typical domestic pressure cooker, the internal pressure is held at about $2atm$ (twice atmospheric pressure) by the safety valve, and the corresponding boiling point of water is approximately $120°C$. Food cooked in this hot water reaches cooking temperatures much faster than it would in an ordinary open pot at $100°C$, because the rate of heat transfer into the food depends on the temperature difference between the water and the food. A cook who can run the water at $120°C$ instead of $100°C$ cuts cooking times by roughly a third for most foods, and by half for tougher meats that require long cooking.

The pressure cooker is therefore not a magic gadget that adds energy to the cooking process. It is a straightforward exploitation of the physics that raising the pressure on a boiling liquid raises its boiling point. The same flame delivers the same heat, but the water reaches a higher temperature before boiling away, and the food consequently cooks faster.

3. Solar water heaters

In Tanzania, where sunshine is abundant and electricity for water heating is expensive, solar water heaters are a common feature on the roofs of homes, hotels, and institutions. A solar water heater consists of two components: a dark absorber panel exposed to the sun, and an insulated storage tank connected to the panel by a pair of pipes. Water flows by natural convection from the tank through the absorber, warming as it passes through, and back into the tank where it remains until a tap is opened below.

Three properties of water make the solar water heater work. **The first** is water's large specific heat capacity. A tank containing $200L$ (and therefore $200kg$) of water can absorb about $840\,000J$ of heat for every degree of temperature rise, which is a substantial buffer against fluctuations in sunlight during the day. **The second** is water's ready availability and low cost; no more expensive working fluid would be justified for domestic use. **The third** is the fact that water flows readily and mixes well, so hot water from the absorber distributes evenly through the tank by convection.

A well-designed solar water heater in Dar es Salaam can warm $200L$ of water from an inlet temperature of $20°C$ to a delivery temperature of $50°C$ or more during a single sunny day, and the insulation of the storage tank keeps most of that heat until the evening when hot water is actually wanted for bathing. The physics of the system is entirely the physics we have already developed: heat absorbed by the panel is used to raise the temperature of the water, the quantity is $Q = mc\Delta T$, and the storage time is set by how well the tank is insulated.

4. Cooling systems for engines and generators

Every internal combustion engine produces more heat than useful work, and the excess heat must be removed continuously if the engine is not to overheat and seize. In almost all modern engines, the heat is removed by circulating water through passages in the engine block, then through a radiator where the heat is released to the outside air. Water is chosen for the same reason it is chosen as a solar-heating fluid: its specific heat capacity is the highest of any common liquid, and a small mass of water flowing at a moderate rate can carry away a very large amount of heat.

In a typical car engine producing $50kW$ of waste heat, a coolant flow rate of about $1kg\,s^{-1}$ is sufficient to keep the engine block at a safe operating temperature, provided the water returns to the engine from the radiator at a temperature about $12K$ below the block temperature. The calculation is direct: $Q/t = \dot{m}c\,\Delta T$, which gives $\dot{m} = 50\,000W/(4200J\,kg^{-1}K^{-1} \times 12K) = 1.0kg\,s^{-1}$ of water. Were the coolant something with a smaller specific heat capacity, a much larger flow rate would be needed, and the pump and radiator would have to be correspondingly larger. Water is therefore not just a convenient coolant; it is the best practical coolant for the job, and its dominance in engine design is a direct consequence of its thermal properties.

Large diesel generators in Tanzanian factories, hospitals, and off-grid communities use the same principle, scaled up. The cooling fluid is usually water plus additives to prevent corrosion and to lower the freezing point, and the radiator is correspondingly larger than a car's, but the physics is identical.

5. Steam as a heat-transfer agent

Every kilogram of steam that condenses on a cold surface releases the latent heat of vapourisation $L_v = 2.26 \times 10^6 J\,kg^{-1}$ directly into that surface. This is far more heat per unit mass than any non-phase-change heat-transfer process can deliver, and it is the reason steam is used as a heat-delivery fluid in a remarkable range of industrial applications.

In hospitals, autoclaves use steam at about $120°C$ to sterilise surgical instruments. The steam penetrates into every cavity of the instruments, condenses on their cold surfaces, and in condensing delivers the latent heat that heats the metal to sterilising temperature. The combination of high temperature and high heat-delivery rate kills bacteria and viruses far more reliably than dry heating, which cannot easily reach the interiors of complex instruments.

In food processing plants, steam is used for cooking, pasteurising, and sterilising because it can be distributed cheaply through pipes to any location where heat is needed and delivers a predictable, controllable amount of heat wherever it condenses. In industrial heating systems for factory floors and large buildings, the same principle applies: steam generated in a central boiler is piped to the areas that need heating, condenses there, and the condensate returns to the boiler to be re-evaporated. The latent heat of vapourisation effectively acts as a thermal capacity reservoir that is orders of magnitude larger than the specific-heat-capacity reservoir of any non-phase-change fluid.

6. Regulation of body temperature

The human body is a thermal system that maintains an internal temperature of about $37°C$ with remarkable precision, despite external conditions that can range from $-20°C$ on a cold mountain to $40°C$ in the tropics. Three mechanisms from this chapter work together to make this regulation possible.

The first is the body's large thermal capacity. An adult human is roughly 60% water by mass, and the specific heat capacity of water is the largest of any common substance, so the body's thermal capacity is comparable to that of a bucket of water. This large thermal capacity buffers the body against rapid temperature swings: a sudden blast of cold or heat takes tens of seconds to change the body's core temperature measurably, giving the body's control systems time to respond.

The second is sweating, which exploits the latent heat of vapourisation of water. A person doing moderate physical work in a hot environment can sweat at a rate of about 1L h^{-1} of water, and the evaporation of each litre of that sweat absorbs about $2.26 \times 10^6\text{J}$ from the skin, cooling the body very efficiently. This is the dominant cooling mechanism at high external temperatures, and it is why dehydration is dangerous in hot climates: a body that cannot replenish its water supply loses its ability to sweat and can overheat within hours.

The third is shivering, which is a metabolic heating mechanism. When the body's core temperature drops, rapid involuntary muscle contractions are triggered, and these contractions convert metabolic energy directly into heat. Shivering can raise the body's heat production by a factor of four or five for short periods, which is usually sufficient to restore core temperature until more permanent measures (warm clothing, shelter) can be found.

The human body is therefore a living thermal system that illustrates nearly every idea in this chapter at once: its temperature is measured by thermometers, its behaviour is stabilised by its specific heat capacity, its cooling uses the latent heat of vapourisation of water, and its heating uses the conversion of metabolic chemical energy into heat.

Theory and you have travelled across six applications of thermal physics and watched every quantity in this chapter appear at least once. Time for thermal equilibrium with two worked examples before miscellaneous worked examples.

BINDER Example 23

A domestic pressure cooker operates at an internal gauge pressure of 1.0atm, so that the absolute pressure inside is $\text{P} = 2.0\text{atm} = 2.03 \times 10^5\text{Pa}$. Under these conditions, the boiling point of water rises to approximately $120°C$. A cook prepares a dish that takes 60min to cook in an open pot of water boiling at $100°C$. As a first approximation, assume that cooking time is roughly inversely proportional to the difference between the water's boiling temperature and room temperature ($25°C$). Calculate (a) the cooking time in the pressure cooker at $120°C$, and (b) the percentage time saved.

Solution

(a) The temperature difference driving the cooking is $\Delta\text{T}_{\text{cook}} = \text{T}_{\text{boil}} - \text{T}_{\text{room}}$. Let t_1 and t_2 be the cooking times in the open pot and the pressure cooker respectively, and let ΔT_1 and ΔT_2 be the corresponding temperature differences.

Open pot: $\Delta\text{T}_1 = 100°C - 25°C = 75°C$, $\text{t}_1 = 60\text{min}$.

Pressure cooker: $\Delta\text{T}_2 = 120°C - 25°C = 95°C$.

Under the first-approximation assumption that cooking time is inversely proportional to ΔT:

$$\frac{\text{t}_2}{\text{t}_1} = \frac{\Delta\text{T}_1}{\Delta\text{T}_2}$$

$$\text{t}_2 = \text{t}_1 \times \frac{\Delta\text{T}_1}{\Delta\text{T}_2} = 60\text{min} \times \frac{75°C}{95°C} = 47.4\text{min}$$

The cooking time in the pressure cooker is roughly 47.4min.

(b) The time saved is $\text{t}_1 - \text{t}_2 = 60\text{min} - 47.4\text{min} = 12.6\text{min}$. As a percentage of the original cooking time:

$$\text{Percentage saved} = \frac{12.6\text{min}}{60\text{min}} \times 100\% = 21\%$$

Under the given approximation, the cook saves about 21% of the cooking time by using the pressure cooker.

Making Sense of the Answer: *The approximation used here treats cooking as a simple heat-transfer process whose rate scales linearly with the temperature difference between the water and room temperature. In reality, cooking involves chemical reactions (denaturation of proteins, gelatinisation of starches, and many others) whose rates depend on temperature in more complicated ways, usually exponentially rather than linearly. The* 21% *figure is therefore an order-of-magnitude estimate rather than an exact prediction; real pressure cookers often save between* 25% *and* 50% *of cooking time depending on the dish, with the larger savings for dishes whose cooking relies on chemical changes that accelerate strongly with temperature.*

Think Like a Physicist: *When a problem gives you a proportionality assumption, check whether the assumption is physically reasonable and what its limits are. Inverse proportionality to a temperature difference is a reasonable first approximation for some processes (simple heat conduction, for example) but a poor one for others (chemical reactions, which typically follow Arrhenius's exponential law). Noting the assumption explicitly and discussing its limits is part of honest physics; silently using a questionable assumption is not.*

HOT Example 24

A solar water heater in Dar es Salaam has an absorber panel of area 2.0m^2, oriented to face the midday sun. On a sunny day, the solar power incident on the panel averages 600Wm^{-2} for 8.0hours between 09:00 and 17:00. The panel absorbs 70% of this incident energy, and the rest is lost by reflection and re-radiation. The absorbed energy is used to warm water from an inlet temperature of 20°C to a delivery temperature of 50°C in a well-insulated storage tank. Take $c_{\text{water}} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and the density of water as 1000kg m^{-3}. Calculate:

(a) the total solar energy incident on the panel during the eight-hour period,

(b) the energy actually absorbed by the panel, and

(c) the maximum volume of water, in litres, that can be heated from 20°C to 50°C using the absorbed energy.

Solution

(a) Total solar energy incident on the panel is the product of the incident power, the panel area, and the duration:

$$E_{\text{incident}} = (\text{power per area}) \times (\text{area}) \times (\text{time})$$

$$E_{\text{incident}} = 600\text{W m}^{-2} \times 2.0\text{m}^2 \times 8.0\text{h} \times 3600\text{s h}^{-1}$$

$$E_{\text{incident}} = 600\text{W m}^{-2} \times 2.0\text{m}^2 \times 2.88 \times 10^4\text{s} = 3.456 \times 10^7\text{J}$$

(b) The panel absorbs 70% of the incident energy:

$$E_{\text{absorbed}} = 0.70 \times E_{\text{incident}} = 0.70 \times 3.456 \times 10^7\text{J} = 2.42 \times 10^7\text{J}$$

(c) The absorbed energy heats water from 20°C to 50°C, a temperature rise of $\Delta T = 30\text{K}$.
Applying $E = mc\Delta T$ and solving for m:

$$m = \frac{E_{\text{absorbed}}}{c_{\text{water}}\,\Delta T} = \frac{2.42 \times 10^7\text{J}}{4200\text{J kg}^{-1}\text{K}^{-1} \times 30\text{K}} = 192\text{kg}$$

Since water has a density of $1000\text{kg m}^{-3} = 1.0\text{kg L}^{-1}$, the volume is

$$V = \frac{m}{\rho} = \frac{192\text{kg}}{1.0\text{kg L}^{-1}} = 192\text{L}$$

The solar water heater can therefore warm about 192L of water from 20°C to 50°C during a single sunny day (which is enough to supply a family of four or five with hot water for bathing and washing).

Making Sense of the Answer: *The* 192L *figure is consistent with the storage capacities of commercial solar water heaters sold in Tanzania, which commonly range from* 100L *for a small household to* 300L *for a larger family. The sizing calculation is straightforward energy balance: solar power in, absorber efficiency, heat required to warm the water, and the answer is the volume. Engineers designing a solar water heater for a particular site run this calculation with local sunshine data and local water demand to decide the panel size and the tank capacity.*

Think Like a Physicist: *A solar-heating sizing calculation is always a three-step operation: total energy incident, efficiency correction to find the absorbed energy, and division by* $c\,\Delta T$ *to find the mass of water that can be heated. Every step uses quantities we have already developed in this chapter. The only additional input the engineer needs is the local solar insolation, which varies with latitude, season, and weather, and which is published in engineering handbooks for every major city on the planet.*

We have now reached the end of Chapter 2. Chapter 2 has handed you a thermometer, a calorimeter, a pressure cooker, and a phase diagram. Before we let you out of the kitchen, the chapter would like to see what you can do with all four at once. The miscellaneous worked examples are next.

MISCELLANEOUS WORKED EXAMPLES ON TEMPERATURE AND HEAT

Example 25

(a) Explain why two identical kettles filled with water at the same initial temperature and placed on identical flames take different times to reach boiling if one of them contains twice as much water as the other.

(b) A 10kg block slides from rest down a rough incline of length 5.0m inclined at 30° to the horizontal. The coefficient of kinetic friction between the block and the incline is 0.40. Assuming that all the heat generated by friction is absorbed by the block itself, and taking the specific heat capacity of the block as $900\text{J kg}^{-1}\text{K}^{-1}$ and $g = 9.81\text{m s}^{-2}$, calculate (i) the heat generated by friction over the full length of the incline, (ii) the temperature rise of the block.

Solution

(a) The heat required to raise the water to boiling is $Q = mc\Delta T$. With the same initial temperature and the same final temperature (boiling point), ΔT and c are the same for both kettles, so Q increases directly with m. The kettle holding twice as much water requires twice as much heat. Since both flames deliver heat at the same rate, the larger kettle takes twice as long to reach boiling.

(b) The normal force on the block is $N = mg\cos\theta$

The friction force is $f = \mu N = \mu mg\cos\theta = 0.40 \times 10\text{kg} \times 9.81\text{ms}^{-2} \times \cos 30^\circ = 33.98\text{N}$.

(i) Heat generated by friction over the full length:

$$Q = fL = 33.98\text{N} \times 5.0\text{m} = 169.9\text{J}$$

(ii) Temperature rise of the block:

$$\Delta T = \frac{Q}{mc} = \frac{169.9\text{J}}{10\text{kg} \times 900\text{J kg}^{-1}\text{K}^{-1}} = 1.89 \times 10^{-2}\text{K}$$

The block's temperature rises by about 0.019K, a tiny change reflecting the large thermal capacity of a 10kg metal block.

Example 26

(a) Explain why a mercury-in-glass thermometer and a platinum resistance thermometer, both carefully calibrated at the ice and steam points, give slightly different readings when placed in the same water bath at $50°\text{C}$, and state which of the two readings can be trusted more for precise work.

(b) A thermocouple has the Seebeck relation $E = a\theta + b\theta^2$ with $a = 40\mu\text{V/°C}$ and $b = -0.040\mu\text{V/°C}^2$, with the cold junction held at $0°\text{C}$. The measured EMF is 9.00mV. Find the two possible hot-junction temperatures consistent with this reading, and state what additional information would resolve the ambiguity.

Solution

(a) The interpolation formula assumes that the thermometric property varies linearly with temperature between the two fixed points, but in reality the expansion of mercury and the resistance of platinum each have a slight curvature. Forced to agree at the ice point and at the steam point, the two thermometers disagree at every intermediate temperature by amounts that reflect the shape of their respective graphs. The platinum resistance reading is generally trusted more for precise scientific work because the resistance of pure platinum has a smaller departure from linearity than the volume of mercury, and because platinum thermometers can be read with an electrical bridge to far higher precision than a mercury column can be read by eye.

(b) Setting $E = 9000\mu\text{V}$ in the Seebeck equation:

$$-0.040\theta^2 + 40\theta = 9000$$
$$0.040\theta^2 - 40\theta + 9000 = 0$$

Solving the quadratic by calculator gives:

$$\theta_1 = 341.9°\text{C} \quad \text{and} \quad \theta_2 = 658.1°\text{C}$$

The ambiguity arises because the EMF-versus-temperature curve is parabolic and peaks at the neutral temperature $\theta_N = -a/(2b) = 500°\text{C}$, which lies between the two roots. The same EMF corresponds to one temperature below the peak and one above. An additional measurement would resolve the ambiguity; for example, a rough check with a separate pyrometer, **or** knowledge of whether the hot junction is being heated up or cooled down (since the EMF changes monotonically on each side of the peak).

Example 27

(a) Explain why coastal cities like Dar es Salaam experience a smaller daily variation in temperature than inland cities like Dodoma, using a single thermal property of water to anchor your answer.

(b) Two thermocouples are connected in series. The first has $a_1 = 40\mu\text{V/°C}$ and $b_1 = -0.030\mu\text{V/°C}^2$ in the form $E = a\theta + b\theta^2$. The second has $a_2 = 60\mu\text{V/°C}$ and $b_2 = -0.045\mu\text{V/°C}^2$ in the same form. The cold junctions are at $0°\text{C}$ and both hot junctions are at $\theta = 200°\text{C}$. Calculate the total EMF measured by a voltmeter connected across the series combination.

Solution

(a) The sea has a much larger thermal capacity than the surrounding land, because water's specific heat capacity ($4200\text{J kg}^{-1}\text{K}^{-1}$) is several times that of sand or rock, and because the sun's heat penetrates many metres down into the sea while only a few centimetres down into the land. The sea therefore warms slowly by day and cools slowly by night, and the air above it tracks the sea's steady temperature. A coastal city, sitting next to this huge thermal reservoir, experiences moderated day-night swings. An inland city, surrounded by land that warms and cools rapidly, experiences larger swings.

(b) The EMF of each thermocouple at $\theta = 200°\text{C}$:

$$E_1 = a_1\theta + b_1\theta^2 = (40\mu\text{V/°C})(200°\text{C}) + \left(-0.030\mu\text{V/°C}^2\right)(200°\text{C})^2 = 6800\mu\text{V}$$
$$E_2 = a_2\theta + b_2\theta^2 = (60\mu\text{V/°C})(200°\text{C}) + \left(-0.045\mu\text{V/°C}^2\right)(200°\text{C})^2 = 10200\mu\text{V}$$

The total EMF in series is $E_{total} = E_1 + E_2$:

$$E_{total} = 6800\mu V + 10200\mu V = 17000\mu V = 17.0mV$$

Example 28

(a) A nurse dips a clinical thermometer in a cold sterilising bath before placing it in a patient's mouth. Explain what physical processes occur between the thermometer and the mouth during the first thirty seconds after the thermometer is inserted, and explain why the nurse must wait before taking the reading.

(b) A laboratory requires thermometers for the following four measurements. For each task, state which type of thermometer (mercury-in-glass, platinum resistance, thermocouple, or optical pyrometer) is most suitable, and give one physical reason supporting your choice: (i) rapidly changing temperatures of exhaust gas leaving a car engine, (ii) the temperature inside a blast furnace at $1600°C$, (iii) the precise temperature of a calibration bath at $50°C$ required to the nearest $0.01°C$, (iv) a child's body temperature at a clinic.

Solution

(a) The thermometer is initially at the temperature of the sterilising bath, which is lower than the mouth's body temperature. When placed in the mouth, heat flows from the mouth (hotter) to the thermometer (colder) by conduction. The thermometer's mercury absorbs this heat and expands, and the mercury column rises steadily up the capillary. The mouth loses a tiny amount of heat to the thermometer in the process, but because the thermometer's thermal capacity is very small compared with the mouth's, the mouth's temperature barely changes. The nurse waits because the thermometer has not yet reached thermal equilibrium with the mouth; a premature reading would give a value below the true body temperature. Only after thermal equilibrium has been reached, when heat flow ceases, does the thermometer's temperature equal the mouth's temperature, and the reading is valid.

(b)

(i) **Thermocouple**, because its small junction has a negligible thermal mass and responds in milliseconds to temperature changes.

(ii) **Optical pyrometer**, because no metal-based thermometer can survive temperatures above about $1500°C$, and a pyrometer reads temperature from radiation emitted by the hot body without making physical contact.

(iii) **Platinum resistance thermometer**, because it provides the highest precision of any standard laboratory thermometer at moderate temperatures, with resolution of $0.01°C$ readily achievable using a Wheatstone bridge.

(iv) **Mercury-in-glass clinical thermometer**, because the range of $35°C$ to $42°C$ is ideal for body temperature and the small mercury bulb delivers adequate precision without disturbing the patient.

Example 29

(a) Explain why the thermodynamic temperature scale does not depend on the thermometric substance used, whereas the Celsius scale defined in the nineteenth century does depend on the substance, and explain why the modern kelvin is defined using a single fixed point rather than two.

(b) A mercury thermometer is faulty: its "$0°C$" mark has been engraved where the mercury actually sits at $5°C$, and its "$100°C$" mark has been engraved where the mercury actually sits at $102°C$. The thermometer is placed in a beaker of water being heated by a $200W$ heater, and reads $75°C$ after the heater has been on for some time. Calculate the true temperature of the water at the moment of reading.

Solution

(a) The thermodynamic scale is built on the ideal-gas relation $P \propto T$ at constant volume, a physical law that every gas obeys exactly in the low-density limit and approximately at ordinary densities. Since the relation is the same for every gas, the scale derived from it is the same regardless of which gas is used. The empirical Celsius scale, by contrast, is built on linear interpolation between two fixed points using the thermometric property of a particular substance (mercury expansion, platinum resistance, or gas pressure), and each of these properties has its own slight non-linearity with temperature. Different substances give slightly different scales between the fixed points.

One fixed point is enough for the thermodynamic scale because the ideal-gas law already fixes the shape of the P-T relation as a straight line through the origin. Only the scale of T remains to be set, and this requires a single reference point. The triple point of water ($T_{tr} = 273.16K$) is used because it is more reproducible than any other accessible physical state.

(b) Applying the faulty-thermometer interpolation formula with $\theta_f = 75°C$, $\theta_{0f} = 5°C$, and $\theta_{100f} = 102°C$:

$$\theta = \left(\frac{\theta_f - \theta_{0f}}{\theta_{100f} - \theta_{0f}}\right) \times 100°C$$

$$\theta = \left(\frac{75°C - 5°C}{102°C - 5°C}\right) \times 100°C = \left(\frac{70°C}{97°C}\right) \times 100°C = 72.16°C$$

The true temperature of the water is approximately $72.2°C$.

Example 30

(a) A drop of water falls on a hot iron at 200°C and is observed to dance across the surface for several seconds before evaporating. Explain in terms of latent heat why evaporation takes so long, and explain why a drop of alcohol would disappear almost instantly under the same conditions.

(b) A laboratory has three thermometers all calibrated at the ice and steam points. A mercury-in-glass thermometer measures column lengths $\ell_0 = 25\text{cm}$, $\ell_{100} = 45\text{cm}$, and $\ell_\theta = 30\text{cm}$ in a room. A platinum resistance thermometer in the same room measures $R_0 = 60\Omega$, $R_{100} = 75\Omega$, $R_\theta = 63.7\Omega$. A thermocouple measures $E_0 = 0\text{mV}$, $E_{100} = 5.5\text{mV}$, $E_\theta = 1.50\text{mV}$. Calculate the room temperature on each of the three scales, and state which value the laboratory should record as the official room temperature, giving a brief reason.

Solution

(a) Even at 200°C the hot iron delivers heat at a finite rate. For each gram of water to evaporate, the iron must supply the latent heat of vapourisation L_v of water, which is a very large quantity of energy per unit mass. The time for the drop to evaporate is proportional to the energy required divided by the rate at which heat is delivered. Alcohol has a substantially smaller latent heat of vapourisation than water, so an alcohol drop of similar size needs much less energy to evaporate completely. With the same hot iron delivering heat at the same rate, the alcohol drop disappears in a much shorter time than the water drop, which is close to instant under the same conditions.

(b) Applying the interpolation formula to each thermometer.
Mercury scale:

$$\theta_{Hg} = \left(\frac{\ell_\theta - \ell_0}{\ell_{100} - \ell_0}\right) \times 100°\text{C} = \left(\frac{30\text{cm} - 25\text{cm}}{45\text{cm} - 25\text{cm}}\right) \times 100°\text{C} = 25.00°\text{C}$$

Platinum scale:

$$\theta_R = \left(\frac{R_\theta - R_0}{R_{100} - R_0}\right) \times 100°\text{C} = \left(\frac{63.7\Omega - 60.0\Omega}{75.0\Omega - 60.0\Omega}\right) \times 100°\text{C} = 24.67°\text{C}$$

Thermocouple scale:

$$\theta_{tc} = \left(\frac{E_\theta - E_0}{E_{100} - E_0}\right) \times 100°\text{C} = \left(\frac{1.50\text{mV} - 0\text{mV}}{5.5\text{mV} - 0\text{mV}}\right) \times 100°\text{C} = 27.27°\text{C}$$

The laboratory should record the platinum resistance thermometer's reading of 24.67°C, because among the three instruments the platinum thermometer is the most accurate and its reading is closest to the true thermodynamic temperature.

Example 31

(a) Explain why a thin aluminium *sufuria* heats up quickly over a flame while a heavy cast-iron pot takes much longer, even though aluminium has a higher specific heat capacity than iron.

(b) In a method-of-mixtures experiment, a copper calorimeter of mass 80g contains 200g of water. The water temperature is monitored by a thermocouple with Seebeck equation $E = a\theta + 1/2\, b\theta^2$ where $a = 50\mu\text{V/°C}$ and $b = -0.040\mu\text{V/°C}^2$, cold junction at 0°C. The thermocouple initially reads $E_1 = 1.20\text{mV}$. A 200g block of an unknown metal is heated to 120°C and dropped into the calorimeter. At thermal equilibrium the thermocouple reads $E_2 = 1.60\text{mV}$. Take $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$, $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$. Calculate the specific heat capacity of the unknown metal.

Solution

(a) The time a pot takes to heat up depends on its thermal capacity $C = mc$, not on its specific heat capacity alone. Aluminium has a smaller density than iron, so for a sufuria of a given size the aluminium version has less mass than the iron version. Although aluminium's specific heat capacity ($900\text{J kg}^{-1}\text{K}^{-1}$) is twice that of iron ($450\text{J kg}^{-1}\text{K}^{-1}$), the mass difference more than compensates, so the product mc is smaller for the aluminium sufuria. A smaller thermal capacity means less heat is needed per degree of temperature rise, so the aluminium sufuria reaches cooking temperature faster.

(b) Converting the two thermocouple readings to temperatures. Setting $E = a\theta + \frac{1}{2}b\theta^2$ and solving for θ.

For $E_1 = 1200\mu\text{V}$:

$$50\theta + \frac{1}{2}(-0.040)\theta^2 = 1200$$

$$0.020\theta^2 - 50\theta + 1200 = 0$$

$$\boldsymbol{\theta = 24.23°\text{C}} \quad (2475.77 \text{ is not practical})$$

Since the calorimeter was at laboratory temperature, the smaller root is practically valid in this case.
The initial temperature of calorimeter, $\theta_1 = 24.23°\text{C}$.
For $E_2 = 1600\mu\text{V}$:

$$0.020\theta^2 - 50\theta + 1600 = 0$$

$$\boldsymbol{\theta = 32.42°\text{C}} \quad (2467.58 \text{ is not practical})$$

Since the metal block started at $120°C$ and is cooling as it loses heat to the calorimeter, the equilibrium temperature θ must be less than $120°C$. So the larger root is a mathematical artefact of the quadratic and has no practical meaning.
The final temperature of calorimeter, $\theta_2 = 32.42°C$.
Applying the heat balance. Heat lost by the metal block = heat gained by water + heat gained by calorimeter:

$$m_M c_M(T_{block} - \theta_2) = (m_w c_w + m_{cal} c_{Cu})(\theta_2 - \theta_1)$$

$$c_M = \frac{(m_w c_w + m_{cal} c_{Cu})(\theta_2 - \theta_1)}{m_M(T_{block} - \theta_2)}$$

$$c_M = \frac{(0.200\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} + 0.080\text{kg} \times 400\text{J kg}^{-1}\text{K}^{-1})(32.42°C - 24.23°C)}{0.200\text{kg} \times (120°C - 32.42°C)} = 407.8\text{J kg}^{-1}\text{K}^{-1}$$

Specific heat capacity of the metal is $407.8\text{J kg}^{-1}\text{K}^{-1}$

Example 32

(a) Explain why a thermometer placed in a beaker of boiling water at standard atmospheric pressure reads $100°C$ continuously for several minutes even though the flame is still supplying heat, and state what happens to the heat during this period.

(b) **Kipute** determines the specific latent heat of fusion of ice by the method of mixtures, using a faulty mercury thermometer that reads $2°C$ at the true ice point and $98°C$ at the true steam point. She places 300g of water in a copper calorimeter of mass 100g; the thermometer initially reads $24°C$. She then adds 20g of ice at the true ice point and stirs until all the ice has melted. The final thermometer reading is $18°C$. Taking $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$, calculate the **correct** value the specific latent heat of fusion of ice from Kipute's data.

Solution

(a) When water is at its boiling point, all the heat supplied by the flame is used to change the phase of the water from liquid to vapour rather than to raise its temperature. During the boiling, the molecules acquire the latent heat of vapourisation L_v, which is used to break the intermolecular forces that hold them in the liquid phase and do work against the atmospheric pressure. The heat therefore appears as an increase in molecular potential energy (and a small amount of external work), not as an increase in molecular kinetic energy. Since temperature is a measure of the average molecular kinetic energy, it stays constant until the last drop of water has been converted to steam.

(b) Converting the thermometer readings to true temperatures using the faulty-thermometer formula.
Initial reading $24°C$:

$$T_1 = \left(\frac{24°C - 2°C}{98°C - 2°C}\right) \times 100°C = 22.92°C$$

Final reading $18°C$:

$$T_2 = \left(\frac{18°C - 2°C}{98°C - 2°C}\right) \times 100°C = 16.67°C$$

Applying the heat balance. Heat lost by the water and calorimeter (cooling from T_1 to T_2) equals heat gained in melting the ice at $0°C$ and then warming the resulting water from $0°C$ to T_2:

$$(m_w c_w + m_{cal} c_{Cu})(T_1 - T_2) = m_i L_f + m_i c_w T_2$$

$$L_f = \frac{(m_w c_w + m_{cal} c_{Cu})(T_1 - T_2) - m_i c_w T_2}{m_i}$$

Where:
$(m_w c_w + m_{cal} c_{Cu})(T_1 - T_2)$
$= (0.300\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} + 0.100\text{kg} \times 400\text{J kg}^{-1}\text{K}^{-1})(22.92°C - 16.67°C) = 8125\text{J}$
$m_i c_w T_2 = 0.020\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times 16.67°C = 1400\text{J}$
Substituting the numerical values:

$$L_f = \frac{8125\text{J} - 1400\text{J}}{0.020\text{kg}} = 3.36 \times 10^5\text{J kg}^{-1}$$

The correct value is $3.36 \times 10^5\text{J kg}^{-1}$.

Example 33

(a) Explain why a wet finger held in the breeze feels cooler than a dry finger, using the concept of latent heat of vapourisation.

(b) A mercury-in-glass thermometer has a mercury bulb of mass 20g (treated as pure mercury, $c_{Hg} = 140\text{J kg}^{-1}\text{K}^{-1}$) and is at room temperature $25°C$. It is dropped into an insulated beaker containing 100g of water at $80°C$. Taking $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and ignoring the heat capacity of the glass, calculate the final equilibrium temperature. Show what this implies about the design of precision thermometers.

Solution

(a) Water on the finger evaporates into the breeze. Each kilogram of water that evaporates draws the latent heat of vapourisation ($L_v = 2.26 \times 10^6 \text{J kg}^{-1}$) from the finger, and this heat loss is what produces the cooling sensation. A dry finger loses heat only by conduction to the slightly cooler air and by a small amount of radiation, neither of which is as efficient as the evaporative process. The breeze increases the rate of evaporation by carrying away the saturated air near the finger, exposing the water to fresh dry air, which speeds up the process and intensifies the cooling.

(b) Applying the heat balance. Heat lost by the water equals heat gained by the thermometer:

$$m_w c_w (T_{water,i} - T_{eq}) = m_t c_t (T_{eq} - T_{therm,i})$$

Substituting values with $T_{water,i} = 80°\text{C}$ and $T_{therm,i} = 25°\text{C}$:

$$0.100\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times (80°\text{C} - T_{eq}) = 0.020\text{kg} \times 140\text{J kg}^{-1}\text{K}^{-1} \times (T_{eq} - 25°\text{C})$$

Solving for T_{eq} gives:

$$T_{eq} = 79.64°\text{C}$$

The water has dropped from 80.00°C to 79.64°C, a change of 0.36°C. The thermometer, in reading this final temperature, has itself disturbed the water it was supposed to measure. This is why precision thermometers are designed with very small bulbs: the smaller the mass (and hence thermal capacity) of the thermometer's working substance, the less it disturbs the body being measured, and the closer its equilibrium reading comes to the body's original temperature. A clinical thermometer with a bulb of only 0.2g instead of 20g would disturb the water by only a hundredth as much.

Example 34

(a) Explain why a pressure cooker cooks food faster than an open pot, using your knowledge of the pressure-dependence of the boiling point of water.

(b) The specific latent heat of vapourisation of water is determined by an experiment in which steam at 100°C is passed into a calorimeter containing water. A copper calorimeter of mass 80g contains 300g of water, and the water temperature is monitored by a platinum resistance thermometer with $R_0 = 50.0\Omega$ at the ice point and $R_{100} = 70.0\Omega$ at the steam point. The platinum thermometer initially reads $R_1 = 53.0\Omega$. Steam is then passed in until 20g of it has condensed, at which point the platinum reads $R_2 = 60.0\Omega$. Taking $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$, calculate the experimental value of L_v obtained from the data.

Solution

(a) The boiling point of water rises with the pressure applied to its surface. In a pressure cooker the steam produced by boiling cannot escape, so it accumulates until the internal pressure has risen well above atmospheric, typically to about 2atm. At this elevated pressure, water boils at a temperature approximately 120°C, twenty degrees higher than in an open pot. The food inside the cooker is therefore surrounded by hot water at 120°C instead of 100°C, and the higher temperature speeds up the chemical reactions that constitute cooking. The cooker does not add energy to the food; it allows the water to reach a higher cooking temperature before boiling away.

(b) Convert the platinum readings to temperatures:

$$T_1 = \left(\frac{R_1 - R_0}{R_{100} - R_0}\right) \times 100°\text{C} = \left(\frac{53.0\Omega - 50.0\Omega}{70.0\Omega - 50.0\Omega}\right) \times 100°\text{C} = 15.0°\text{C}$$

$$T_2 = \left(\frac{R_2 - R_0}{R_{100} - R_0}\right) \times 100°\text{C} = \left(\frac{60.0\Omega - 50.0\Omega}{70.0\Omega - 50.0\Omega}\right) \times 100°\text{C} = 50.0°\text{C}$$

Applying the heat balance. Heat released by steam (condensing at 100°C and then cooling as liquid water from 100°C to T_2) equals heat absorbed by the original water plus the calorimeter (warming from T_1 to T_2):

$$m_s L_v + m_s c_w (100°\text{C} - T_2) = (m_w c_w + m_{cal} c_{Cu})(T_2 - T_1)$$

Substituting the numerical values:

$$0.020\text{kg} \times L_v + 0.020\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times (100°\text{C} - 50.0°\text{C})$$
$$= (0.300\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} + 0.080\text{kg} \times 400\text{J kg}^{-1}\text{K}^{-1})(50.0°\text{C} - 15.0°\text{C})$$

Solving for L_v gives:

$$L_v = 2.051 \times 10^6 \text{J kg}^{-1}$$

Ten examples down. You have cooked food in pressure cookers, dropped iron into ice water, read three thermometers in the same room, and determined the latent heat of ice with a faulty thermometer that the physics quietly sneaked past. The heavy lifting is complete. What remains is Digging Deeper, where the questions keep coming but the physics stops teaching and starts testing.

DIGGING DEEPER EXERCISE 2

EXERCISE 2A: BINDER QUESTIONS

Question 1

Kipanga says: *"If body A is in thermal equilibrium with body B, and body B is in thermal equilibrium with body C, then A is obviously in thermal equilibrium with C. This is just common sense. I don't see why we need to call it the Zeroth Law."*

Explain what the Zeroth Law guarantees that is not obvious from everyday reasoning, and state why this guarantee is needed before any thermometer can be trusted.

Question 2

A student claims: *"A bucket of water at* $20°\text{C}$ *contains more heat than a cup of water at* $80°\text{C}$*, because the bucket has more molecules."* Identify and correct the mistake in this statement.

Question 3

(a) Distinguish between a **thermometric property** and an ordinary physical property of a substance.

(b) Give two examples of physical properties that cannot be used as thermometric properties, and state for each why not.

Question 4

The modern kelvin is defined using a single fixed point (the triple point of water) rather than the two fixed points (ice point and steam point) used by the nineteenth-century Celsius scale. State two distinct disadvantages of the ice point that make it unsuitable as the primary reference in the modern definition.

Question 5

A student dips a thermometer in a bath and waits for the reading to stabilise before recording it.

(a) State what physical condition has been reached when the reading stops changing.

(b) Explain why this condition must be reached before the reading can be trusted.

Question 6

(a) State the property of water, together with its numerical value, that makes water the preferred coolant in car engines and power stations.

(b) Oil has a specific heat capacity of about $2000\text{J kg}^{-1}\text{K}^{-1}$. Explain why oil is not used as a coolant when water is available.

Question 7

Explain, using molecular reasoning, why heat is absorbed when ice melts but released when water freezes, even though both processes occur at exactly $0°\text{C}$.

Question 8

(a) State three distinct differences between evaporation and boiling.

(b) State one thing that the two processes have in common.

Question 9

Explain why precision thermometers are designed with very small bulbs, and state what would happen to the accuracy of a measurement if the bulb's thermal capacity were comparable to that of the body being measured.

EXERCISE 2B: REAL QUESTIONS

Question 10

Kipanga and Kipute are arguing on a hot daladala ride from Chalinze to Dar es Salaam. **Kipanga**, sweating heavily, says: *"It is hotter today than yesterday. I know this because I am sweating more than I did yesterday."* **Kipute** replies: *"Sweating tells you about your body, not about the air. Yesterday I was eating ugali that raised my body heat, and I sweated just as much as you are sweating now."*

(a) State which speaker is using temperature correctly as a physical quantity, and why the other is not.

(b) Explain what a thermometer measures that makes it more reliable than human sweating as a temperature indicator.

Question 11

Nurse Maria at a Bagamoyo clinic is measuring the temperatures of ten children queued for vaccination. She dips her clinical thermometer in cool antiseptic between patients, removes it, shakes it hard to reset the mercury column, and immediately places it in the next child's mouth. A doctor observing her procedure says the readings will be systematically too low.

(a) State two physical reasons why the nurse's procedure produces systematically low readings.

(b) State two changes she could make that would give more reliable readings.

Question 12

A family in Kibaha is deciding between a charcoal stove and a gas cooker for daily cooking. The mother notes that a pot of beans takes about 90min on charcoal and 45min on gas. The father argues: *"The gas flame must be twice as hot as the charcoal flame."* Explain why the father's reasoning is wrong, and state the physical property of the two cookers that actually differs.

Question 13

Mama Juma lives in the cool highlands near Mufindi and uses a hot water bottle to keep her bed warm at night. A salesperson offers her a new bottle of the same size filled with hot cooking oil instead of hot water, claiming "oil stays hot longer". Taking $c_{water} = 4200\,\text{J kg}^{-1}\text{K}^{-1}$ and $c_{oil} = 2000\,\text{J kg}^{-1}\text{K}^{-1}$, and assuming both bottles cool at the same rate of heat loss to the surroundings, advise Mama Juma which bottle will keep her warmer for longer, and explain the physics behind your advice.

Question 14

A long-distance bus driver notices his engine is overheating on a climb to Kitonga. He considers replacing the water in the cooling system with used engine oil, which he has available cheaply.

(a) Explain why this would be a bad idea, using specific heat capacity as the centrepiece of your answer.

(b) State one property other than specific heat capacity that matters for an engine coolant.

Question 15

A Form Five physics student is performing a calorimetry experiment and notices that three thermometers in the laboratory give slightly different readings when placed in the same beaker of warm water: a mercury thermometer reads 48.0°C, a platinum resistance thermometer reads 48.5°C, and a thermocouple reads 47.2°C. The lab technician tells her to *"just use whichever one you have calibrated most recently."*

(a) Explain whether the technician's advice is physically sensible.

(b) State which of the three readings the student should record as the best estimate of the true temperature, and justify your choice.

Question 16

A sambusa vendor in Kariakoo fries her sambusas in a large pot of hot oil. On a quiet day she fries five sambusas at a time. On a busy day she drops twenty sambusas into the same pot at once, and she observes that the oil temperature drops noticeably after adding the sambusas, with the first batch taking longer to cook than usual. Explain this observation using the concept of thermal capacity, and recommend one practical change she could make to avoid the problem.

Question 17

Kipanga wants to cool a bottle of soda quickly. He has two options: place the bottle in a basin of iced water at 0°C, or place it in a basin of water that has been cooled to 4°C and contains no ice. Kipute claims the iced water will cool the soda faster even though the water in both basins is already very cold. Explain why Kipute is right.

Question 18

A cook in a hurry leaves the lid off her pot of boiling stew, saying: *"The flame is the same either way, so the time to boil will be the same whether or not I use the lid."* Explain why the cook is wrong.

Question 19

A domestic pressure cooker is designed with a safety valve that releases steam when the internal pressure exceeds a set limit. A student observes that the valve stays silent for the first ten minutes after the cooker is placed on a strong flame, then begins to hiss and continues to hiss until the flame is switched off.

(a) Explain why the valve stays silent for the first ten minutes and then begins to hiss.

(b) Explain what the continued hissing tells us about the rate of heat input compared with the rate of steam loss.

(c) State whether the cooker's internal temperature drops if the user reduces the flame to a lower setting after the valve has begun to hiss, provided the valve continues to hiss.

EXERCISE 2C: HOT QUESTIONS

Question 20

Water falls from a 200m-high waterfall in the Southern Highlands. Assuming all gravitational potential energy is converted to heat absorbed by the water itself, calculate the temperature rise of the water at the bottom of the fall. Take $g = 9.81\,\text{m s}^{-2}$ and $c_{water} = 4200\,\text{J kg}^{-1}\text{K}^{-1}$. State two physical reasons why the actual temperature rise in a real waterfall is smaller than this calculation predicts.

Question 21

A 500g iron block at 200°C is dropped into a copper calorimeter of mass 100g that contains 300g of water and 50g of ice, all initially at 0°C. Take $c_{Fe} = 450\text{J kg}^{-1}\text{K}^{-1}$, $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$, $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$, and $L_f = 3.34 \times 10^5\text{J kg}^{-1}$. Calculate the final equilibrium temperature of the system.

Question 22

A copper block of mass 800g at 250°C is dropped into an open copper calorimeter of mass 100g that contains 150g of water at 100°C. Take $c_{Cu} = 400\text{J kg}^{-1}\text{K}^{-1}$, $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$, and $L_v = 2.26 \times 10^6\text{J kg}^{-1}$. Calculate the mass of water that boils away before the system reaches equilibrium, and state the final temperature of the contents.

Question 23

A thermocouple is calibrated by placing its hot junction in melting ice (0°C, EMF = 0), boiling water at standard atmospheric pressure (100°C, EMF = 3.800mV), and steam at higher pressure (120°C, EMF = 4.512mV). Assume the Seebeck relation has the form $E = a\theta + 1/2\, b\theta^2$ with the cold junction at 0°C. Determine the constants a and b, and predict the EMF when the hot junction is at 50°C.

Question 24

A shallow pond in the Southern Highlands has an ice layer on its surface. The air above the ice is at -8°C and the water beneath the ice is at 0°C. Take the thermal conductivity of ice as $k = 2.2\text{W m}^{-1}\text{K}^{-1}$, the density of ice as $\rho = 920\text{kg m}^{-3}$, and the latent heat of fusion as $L_f = 3.34 \times 10^5\text{J kg}^{-1}$. Calculate the time required for the ice layer to grow in thickness from 6.0cm to 7.0cm, assuming the air and water temperatures remain constant during this period.

Question 25

A solar water heater in Iringa has an absorber panel of area 2.5m^2 and heats a tank containing 180L of water. On a partly cloudy day, the solar power incident on the panel averages 350W m^{-2} for 9.0hours, and the panel absorbs 65% of the incident energy. The water in the tank starts at 22°C. Take $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and the density of water as 1000kg m^{-3}. Calculate the final temperature of the water in the tank.

Question 26

A platinum resistance thermometer is supposed to give $R_0 = 50.00\Omega$ at the ice point and $R_{100} = 65.00\Omega$ at the steam point, but because of calibration errors it actually gives $R_0 = 50.50\Omega$ at the true ice point and $R_{100} = 64.80\Omega$ at the true steam point. The same thermometer is placed in a bath and reads 58.00Ω.

(a) Calculate the temperature the student would report if she assumed the thermometer was properly calibrated.

(b) Calculate the true temperature of the bath.

(c) Express the difference between (a) and (b) as a percentage of the true temperature.

Question 27

Steam at 100°C is passed into an insulated tank containing 10.0L of water at 20°C, and the steam condenses fully as it enters. Taking $L_v = 2.26 \times 10^6\text{J kg}^{-1}$, $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$, and the density of water as 1000kg m^{-3}, calculate the mass of steam that must condense to bring the water in the tank to 60°C.

Question 28

A pressure cooker holds 2.0L of water that is heated from room temperature 25°C. In an open pot, the water needs only to reach 100°C to begin cooking; in the pressure cooker, the operating temperature is 120°C. Take $c_{water} = 4200\text{J kg}^{-1}\text{K}^{-1}$ and the density of water as 1000kg m^{-3}. The flame supplies heat at a steady rate of 2.0kW.

(a) Calculate the additional heat required to bring the cooker to its operating temperature, compared with simply reaching ordinary boiling.

(b) Calculate the additional time the flame must operate to deliver this additional heat.

Question 29

A wet-and-dry bulb hygrometer has two identical mercury thermometers placed side by side. The dry bulb reads air temperature 30°C; the wet bulb, whose mercury is wrapped in a damp cloth, reads 24°C in still air. The wet bulb is in a steady state in which the rate of heat lost by evaporation of water from the cloth is equal to the rate of heat gained from the surrounding air, which is 0.50W. Taking $L_v = 2.26 \times 10^6\text{J kg}^{-1}$, calculate the rate at which water must be evaporating from the cloth, expressed in grams per hour.

Question 30

A 100g sample of an unknown solid is heated steadily by a 50W electrical heater. The temperature-versus-time graph shows three segments: from $t = 0$ to $t = 60$s the temperature rises linearly from -60°C to -20°C; from $t = 60$s to $t = 140$s the

temperature stays constant at $-20°C$; from $t = 140s$ to $t = 200s$ the temperature rises linearly from $-20°C$ to $50°C$. Calculate the specific heat capacity of the solid phase, the latent heat of fusion at $-20°C$, and the specific heat capacity of the liquid phase.

Question 31

A lead bullet of mass $5.0g$ moving at $350m\ s^{-1}$ strikes and embeds in a wooden block. Assume that all the bullet's kinetic energy is converted to heat absorbed by the bullet alone (the block absorbs none). The bullet's initial temperature is $25°C$. Take the specific heat capacity of lead as $c_{Pb} = 128J\ kg^{-1}K^{-1}$, the melting point of lead as $327°C$, and the latent heat of fusion of lead as $L_{f,Pb} = 2.5 \times 10^4 J\ kg^{-1}$. Determine whether the bullet melts on impact, and if so, calculate the mass of lead that melts.

Question 32

A $50g$ ice cube at $0°C$ is dropped from a height of $100m$ into a still pond at $0°C$. Take $g = 9.81m\ s^{-2}$ and $L_f = 3.34 \times 10^5 J\ kg^{-1}$. Assuming all the gravitational potential energy is converted to heat absorbed by the ice cube alone, calculate the mass of ice that melts on impact, and express it as a percentage of the cube's original mass.

Chapter 3

HEAT TRANSFER

INTRODUCTION

The school kitchen at Miono Secondary had two rules. The first was that no student was allowed inside without supervision. The second was that no student was allowed to touch the cooking equipment, because Mama Salha, the school cook, had once caught a Form Five student tasting her bean stew, and the public lecture that followed had been so fierce that the boy avoided the kitchen forever.

None of this applied when Mr. Akilikubwa was involved. He had been teaching physics at Miono for twenty-three years, and the entire school simply trusted him with things. If he wanted to borrow the kitchen for a lesson, Mama Salha would sigh, wipe her hands on her apron, and say, *"Take it, but if that boy with the big voice breaks my pot, you are buying the replacement."* The boy with the big voice was, of course, Kipanga.

On this Wednesday morning, Mr. Akilikubwa had arrived before the first bell. Sitting on one of the kitchen's large commercial gas burners was an aluminium *sufuria* of water that had been boiling hard for ten minutes, sending a column of steam toward the corrugated iron ceiling. Inside the *sufuria*, resting against the rim with their handles sticking out above the waterline, were two utensils. On the left: a stainless steel cooking spoon, the kind with a long metal handle that runs all the way from the bowl to the tip without interruption. On the right: a wooden *mwiko*, with a thick round handle Mama Salha had been stirring *ugali* with since before any of these students were born.

Both had been sitting in the same boiling water, in the same pot, on the same gas burner, for the same ten minutes. This detail is important, and Mr. Akilikubwa would make sure it was repeated until nobody could pretend they had missed it.

The Form Five physics class filed in through the side door. The kitchen smelled of steam and yesterday's bean stew. They arranged themselves in a semicircle around the burner. Kipanga, who always stood at the front on the grounds that education worked best at close range, positioned himself directly between the two utensil handles. Kipute took a spot two rows back, where she had a clear view of both handles and, more importantly, a clear view of Kipanga.

Mr. Akilikubwa let the silence build. He was excellent at this. Other teachers filled silence with words; Mr. Akilikubwa filled it with expectation, and the expectation always paid a higher dividend.

"Two utensils," he said, when every eye was on him. *"One steel. One wood. Both have been in the same boiling water for the same ten minutes. The same temperature. The same heat available. The same everything. I would like two volunteers. Each will pick up one utensil by the handle, hold it for three seconds, then put it down. Who would like to go first?"*

Kipanga's hand went up with the speed and enthusiasm of a man who had learned absolutely nothing from the previous two chapters.

"Choose your utensil, Kipanga."

Kipanga did not hesitate. He did not pause to ask himself why a physics teacher might be offering this choice in front of thirty witnesses. He simply reached for the stainless steel spoon, because it was shinier, because it looked more professional, and because Kipanga, in the long and colourful history of his physics education, had never once made a decision by thinking about it first.

He grabbed the steel handle.

For approximately one quarter of a second, nothing happened. This was the quarter-second in which the thermal energy stored in the steel, conducted at roughly 50 watts per metre per kelvin, was still travelling from his palm to his pain receptors. If physics had a sense of humour, this was the pause before the punchline.

Then the quarter-second ended.

Kipanga released the spoon with a sound that the English language does not have an adequate word for, but which the Swahili language would describe as something between a yelp, a prayer, and a very short autobiography of regret. The spoon hit the rim of the *sufuria*, bounced once with a bright metallic ring, and clattered onto the concrete floor. Kipanga shook his hand in the air with the frantic energy of a man trying to detach his own fingers from his body, then stared at his palm as though expecting to find a written explanation branded into the skin.

"That," he said, with the wounded dignity of a man who feels the universe owes him an apology, *"is not a fair experiment."*

Mr. Akilikubwa did not respond. He simply turned to the rest of the class. *"Volunteer number two?"*

Kipute walked forward without raising her hand, because Kipute did not compete for things she had already worked out. She picked up the wooden *mwiko* by its handle. She held it for three seconds. She held it for ten. She could have held it all morning. The handle was warm, the way a cup of tea is warm, but it was not hot. It was not painful. It was not even uncomfortable. It had been sitting in exactly the same boiling water, at exactly the same temperature, for exactly the same ten minutes as the

steel spoon that had just reorganised Kipanga's entire relationship with kitchen utensils. She placed the *mwiko* back in the pot and returned to her spot without a word.

Kipanga, still holding his right hand with his left as though it were an injured bird, stared at her with the expression of a man watching someone walk calmly through a door that had just electrocuted him.

"How," he said, in a voice quiet enough to confirm that he was genuinely confused rather than performing confusion for the audience. *"Same water. Same temperature. Same time. The only difference between that spoon and that mwiko is the material. And you are telling me the material is the reason I just lost all feeling in my right hand?"*

"Yes," said Kipute. *"Remember the metal gate two weeks ago? You touched it at 14 degrees and it felt freezing, then sat on the wooden bench at the same 14 degrees and it felt warm. I told you then that your hand was not measuring temperature. It was measuring the rate at which heat flowed out of it. The steel spoon has the same story. It is not hotter than the mwiko. It conducts heat into your skin faster."*

Kipanga opened his mouth. He closed it. He looked at his palm, where a faint pink stripe was forming where the steel handle had pressed. He looked at the *mwiko*, still resting peacefully in the pot. And then he said something nobody in the Form Five physics class had ever heard Kipanga say before.

"That actually makes sense."

Three students in the back row applauded. Kipanga ignored them with considerable effort.

Mr. Akilikubwa stepped forward, picked the steel spoon off the floor, wiped it, and placed it back in the *sufuria*. *"Kipute is exactly right. What your hand felt was not temperature. It was the rate at which thermal energy flowed from the spoon into your skin. Steel has a high thermal conductivity; it delivers heat fast. Wood has a very low thermal conductivity; it delivers the same available heat so slowly your hand barely registers the transfer. The temperature was identical. What was completely different was the speed of delivery. And that single idea, the rate of heat transfer, is the subject of this entire chapter."*

He turned to the small portable whiteboard propped against the tiles and wrote three words in a vertical column:

CONDUCTION CONVECTION RADIATION

"Chapter 2 asked how much heat. How many joules to raise water by one degree, to melt ice, to boil a kettle. All those numbers lived inside $Q = mc\Delta T$ and $Q = mL$."

*"Chapter 3 asks a different question. It asks how fast that heat moves and which road it travels. Does it travel through the material itself, atom by atom, without the material moving? That is **conduction**, and it is the reason Kipanga's hand is still stinging. Does it travel because the material itself moves, carrying heat the way a daladala carries passengers from Mwenge to Mbagala? That is **convection**, and it is the reason the ceiling of your classroom is warmer than the floor. Or does it travel as electromagnetic radiation, needing no material at all, crossing the vacuum of space from the Sun to your face? That is **radiation**, and it is the reason the Earth is not a frozen rock, even though it is surrounded by the coldest emptiness in the universe."*

"Every thermal phenomenon you will meet in the next seventy pages is one of these three mechanisms doing its work, or being deliberately stopped from doing its work. By the end of this chapter, you will know why a blanket keeps you warm in July and why the same blanket wrapped around a block of ice slows its melting. Why birds puff out their feathers and animals curl into a ball. Why ice on a pond grows more slowly as it gets thicker. How a forensic scientist estimates time of death by measuring a body's temperature. And why Kipanga's hand hurts and Kipute's does not."

Kipanga held up his palm for the class to see, turning it slowly so that even the back row could appreciate the evidence.

"Scientific contribution," he announced, with the dignity of a man who has decided that his suffering has advanced human knowledge and therefore was not entirely in vain.

The class laughed. Kipute closed her notebook. And just like that, Chapter 3 had begun.

Chapter 2 showed how much heat is needed to change temperature ($Q = mc\Delta T$) and how much is needed for a phase change ($Q = mL$). It established the energy bookkeeping of thermal physics: if you know the mass, the material, and the temperature change, you can calculate the joules. But it deliberately left one question unanswered. When Kipanga touched the metal gate at 14 °C and it felt freezing, while the wooden bench at the same 14 °C felt warm, Chapter 2 said only that the explanation involved the rate of heat flow and that the full answer would come later. This is later.

Chapter 3 supplies the missing physics. Fourier's law of conduction, derived properly in the conduction section, gives the rate of heat flow through a solid in terms of the material's thermal conductivity, the cross-sectional area, and the temperature gradient. This is the same equation that was borrowed as an unproved tool in Chapter 2's ice pond calculation; it is now derived from first principles and explored in full generality, including radial heat flow through spheres and cylinders. Newton's law of cooling, developed later in this chapter, provides the exponential temperature decay of a body losing heat to its surroundings, the same mathematical form that will reappear in Chapter 11 when we model the discharge of a capacitor through a resistor, because the underlying differential equation is identical. The Stefan-Boltzmann radiation law connects forward to Chapter 5, where the distinction between an isothermal process and an adiabatic process depends entirely on

knowing how quickly a system exchanges heat with its surroundings. If the exchange is fast (good conductor, large surface area, long time available), the process is isothermal. If the exchange is slow (poor conductor, small surface, rapid compression), the process is adiabatic. That judgement requires the tools of this chapter.

THERMAL CONDUCTION

Of the three heat-transfer mechanisms introduced in our kitchen, **conduction** is the one that taught Kipanga his lesson. The steel spoon and the wooden *mwiko* were sitting in the same boiling water at the same temperature, yet only one of them sent a stripe across his palm. The water did not move from the pot to his hand. The spoon did not move. The handle simply stood there in the air above the rim, and somehow the heat travelled along it from the boiling tip to the cool grip in less than a second. That mysterious through-the-metal journey is what physicists call conduction, and the next several pages will turn it from a mystery into a calculation.

Conduction *is the process in which heat flows from the hotter region of a material to the cooler region without any net movement of the material itself.* Picture the steel spoon at the molecular level. Down at the bowl end, the molecules are immersed in 100-degree water and are vibrating with great energy. Up at the handle end, the molecules are surrounded by cool morning air and are vibrating gently. Each vibrating molecule jostles its immediate neighbours. The energetic ones at the bottom shake their neighbours harder; those neighbours shake the next neighbours; and the disturbance ripples up the spoon as a wave of increasingly active vibration. Heat has moved from the bowl to the handle, but no atom of steel has actually travelled with it. The molecules vibrate in place. They simply pass the energy along, like a row of seated people clapping a rhythm down the line.

This passing-along is what makes conduction so different from the other two mechanisms. In convection, the hot fluid itself rises and physically carries the energy with it. In radiation, no medium is needed at all. But in conduction, the material stays where it is, and the energy migrates through it from atom to atom. This explains immediately why **solids conduct heat best** (*their atoms are packed close together and pass vibrations efficiently*), **liquids conduct moderately** (*their molecules are close but can shift around, dissipating the wave*), and **gases conduct very poorly** (*their molecules are so far apart that vibrations rarely reach a neighbour at all*).

Fourier's Law of Heat Conduction

Saying that heat flows from hot to cold is true but useless. To do physics with conduction, we need a number. We need to know exactly *how many joules of energy will travel through a slab of material every second, given how thick the slab is, how big its faces are, what it is made of,* and *how big the temperature difference is between its two faces.* The relationship that gives us this number is called **Fourier's law**, and we will build it from observations rather than memorise it from a book.

Imagine a uniform slab of material with **cross-sectional area A** and **thickness L**. One face is held at the **higher temperature θ_1** and the opposite face is held at the **lower temperature θ_2.** After waiting long enough for the slab to settle into a **steady state** (this means *waiting until the temperature at every internal point has stopped changing with time*), heat flows steadily through the slab from the hot face to the cold face. The geometry is shown in the following figure.

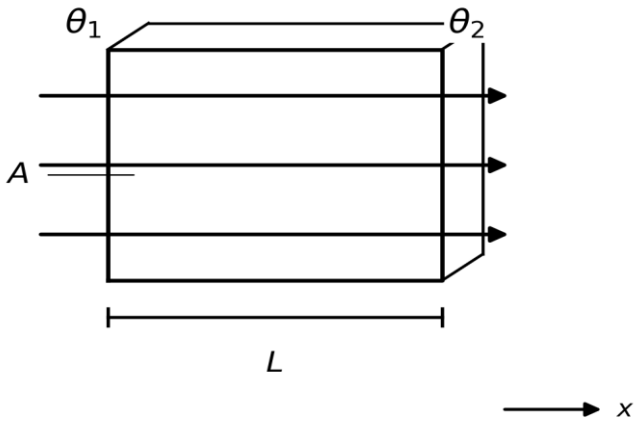


Figure: ***Heat conduction through a uniform slab of cross-sectional area A and thickness L.*** *The hot face on the left is at temperature θ_1 and the cold face on the right is at θ_2. Heat flows in the +x direction, from hot to cold.*

Let dQ be the small quantity of heat that crosses the slab in a small time dt. Experiment shows that dQ depends on four factors, each in a way that should match your intuition.

First, dQ is directly proportional to the area A of the slab. A bigger glass window loses more heat than a smaller one of the same thickness, just as more lanes on a highway can carry more cars per minute than fewer lanes.

Second, dQ is directly proportional to the time dt for which the heat is allowed to flow. Twice as long, twice as much heat. This is so obvious it almost does not need stating, but Fourier was a careful man and did state it.

Third, dQ is directly proportional to the temperature difference $(\theta_1 - \theta_2)$ between the two faces. Doubling the temperature difference doubles the rate at which heat flows. This is why a window leaks more heat on a freezing night than on a mild evening.

Fourth, dQ is inversely proportional to the thickness L of the slab. A thick wall conducts heat more slowly than a thin one, which is why a 30 cm brick wall keeps a house cooler in midday Dar es Salaam than a 5 mm sheet of glass would.

Combining all four observations into a single proportionality:

$$dQ \propto \frac{A(\theta_1 - \theta_2)\, dt}{L}$$

To turn the proportionality into an equality, we introduce a constant of proportionality. This constant cannot be universal, because different materials clearly conduct heat at very different rates: a copper bar feels almost as hot as the boiling water at its far end within seconds, whereas a wooden stick of the same dimensions in the same water remains almost cool. The constant must therefore be a property of the specific material being used. We call it the coefficient of thermal conductivity and give it the symbol k. With this constant inserted:

$$dQ = \frac{k\, A\, (\theta_1 - \theta_2)\, dt}{L}$$

Dividing both sides by dt gives the rate of heat flow, which is exactly the quantity we want:

$$\boldsymbol{\frac{dQ}{dt} = \frac{k\, A\, (\theta_1 - \theta_2)}{L}}$$

This is **Fourier's law** for steady-state conduction through a uniform slab. The left-hand side is the rate at which heat crosses the slab, measured in joules per second, which is watts. The right-hand side is the product of four things you can measure: the thermal conductivity of the material, the area of the slab, the temperature difference between the faces, and the thickness of the slab. Once you know these four numbers, the rate of heat flow follows immediately.

In words the law can be stated as: *the rate at which heat is conducted through a uniform slab in steady state is directly proportional to the cross-sectional area of the slab and to the temperature difference between its two faces, and inversely proportional to the thickness of the slab.*

The Temperature Gradient and the Negative Sign

The form of Fourier's law we have just written is the one that students use most often, but it hides a small subtlety that becomes important the moment we move beyond a uniform slab. Heat does not really care about the temperature difference between the two faces. Heat cares about the rate at which temperature changes with distance, which is to say it cares about the slope of the temperature curve at the point where the heat is passing through.

To see this clearly, imagine cutting a thin slice of thickness dx out of the slab at some depth x, where the temperature on the hot side of the slice is θ and the temperature on the cold side is $(\theta + d\theta)$. Note that $d\theta$ here is a negative quantity, because temperature decreases as we move from hot to cold along the +x direction. The temperature change across the thin slice is $d\theta$, the thickness is dx, and the rate of change of temperature with distance, called the temperature gradient, is $d\theta/dx$. Applying Fourier's law to this thin slice:

$$\boldsymbol{\frac{dQ}{dt} = -kA\,\frac{d\theta}{dx}}$$

The minus sign appears for a precise physical reason that is worth stating slowly. We have chosen the +x direction to point from the hot face toward the cold face, so heat flows in the +x direction; we want dQ/dt to be a positive number when heat is flowing the way we have set up our axes. But $d\theta/dx$ is negative, because temperature is dropping as x increases. Multiplying by a minus sign converts the negative gradient into a positive heat flow rate. The minus sign is therefore not a mathematical inconvenience. It is the equation's way of telling you that heat always flows **down** the temperature gradient, from high to low. If you ever find yourself having computed a **negative** rate of heat flow in a problem, the equation is not broken; it is simply telling you that the heat is actually flowing in the opposite direction to the one you assumed.

In a uniform slab at steady state, the temperature falls linearly from θ_1 at one face to θ_2 at the other, so the gradient $d\theta/dx$ is constant everywhere inside the slab and equals $-(\theta_1 - \theta_2)/L$. Substituting this into the Fourier law immediately recovers our earlier formula, with the minus signs cancelling neatly. The two forms of the law are therefore the same physics, written for two different situations: the gradient form is general and applies even when the temperature varies non-linearly with position (such as inside a sphere or a cylinder, which we will meet later in this section), while the slab form is the special case for uniform geometry.

The Thermal Conductivity, k

The constant k that appeared in our derivation is the most important number in this chapter. It is the coefficient of thermal conductivity, and it is the single number that distinguishes a copper saucepan from a clay one, a glass window from a brick

wall, and Kipanga's steel spoon from Kipute's wooden *mwiko*. Before we tabulate values for different materials, let us extract a careful definition of k from the equation it lives in.

Rearranging Fourier's law to make k the subject:

$$k = \frac{\frac{dQ}{dt}}{\frac{A(\theta_1 - \theta_2)}{L}}$$

Look at what this equation is telling us. The **thermal conductivity of a material** *is the rate of heat flow per unit area per unit temperature gradient.* Equivalently, k *is the quantity of heat that flows per second through a slab of unit thickness and unit cross-sectional area when the temperature difference between its two faces is one kelvin.* A material with a large k delivers a lot of heat per second under modest conditions; a material with a small k delivers very little. The unit of k is watts per metre per kelvin (W m^{-1} K^{-1}), which can also be written as joules per second per metre per kelvin (J s^{-1} m^{-1} K^{-1}). The two are identical, since one watt is one joule per second.

To a very good approximation, the thermal conductivity of any material is independent of temperature over the range of temperatures we will meet in this book. This is a useful simplification: we can treat k as a fixed number for each material, look it up in a table, and use it without worrying that the value might shift midway through a calculation. (At very high temperatures, or for some exotic materials, k does drift; A-level physics ignores this.)

A table of thermal conductivities

The following values cover almost every material you will meet in this chapter and the next. Notice the enormous range, more than four orders of magnitude from silver at the top to air at the bottom.

Material	k / W m^{-1} K^{-1}	Conductor type
Silver	420	excellent conductor
Copper	380	excellent conductor
Aluminium	240	very good conductor
Brass	109	good conductor
Iron	80	good conductor
Lead	35	moderate conductor
Concrete	1.3	poor conductor
Glass	1.0	poor conductor
Brick	0.7	insulator (mild)
Water	0.6	insulator
Wood (*mwiko*)	0.15	insulator
Felt	0.04	insulator (good)
Air (still)	0.026	insulator (excellent)

A few patterns leap out of this table. Metals occupy the top half (silver, copper, aluminium, brass, iron, lead) with conductivities in the tens to hundreds. Non-metallic solids and liquids occupy the middle (concrete, glass, brick, water, wood) with conductivities below 2. Trapped gases sit at the bottom (felt and still air) with conductivities below 0.05. Felt and still air are not chemically magical: they are just gases that have been trapped in a way that prevents them from convecting. We will return to this trapped-air idea repeatedly in this chapter, because it explains everything from blankets to bird feathers to the design of vacuum flasks.

One entry deserves a second look. Stainless steel does not appear in the table; its thermal conductivity is roughly 50 W m^{-1} K^{-1} (the value we quoted in the kitchen as Kipanga's palm was being introduced to physics). Wood is roughly 0.15. The ratio of these two numbers is more than 300, which is why his hand experienced the steel spoon and the *mwiko* as utterly different objects even though they were sitting in the same boiling water. The temperature was identical. The conductivity was 300 times apart. His skin measured the conductivity, not the temperature.

Three short worked examples follow, each chosen to drill a different muscle. None of them will hurt as much as the steel spoon. We promise.

BINDER Example 1

A glass window in a school laboratory has thickness 8.0mm and area 2.0m². The inner surface of the glass is at 20°C and the outer surface is at 0°C. The thermal conductivity of glass is $1.0\text{Wm}^{-1}\text{K}^{-1}$. Calculate the rate of heat loss through the window.

Solution

Using Fourier's law:

$$\frac{dQ}{dt} = \frac{k\,A\,(\theta_1 - \theta_2)}{L}$$

Substituting the given values:

$$\frac{dQ}{dt} = \frac{(1.0\text{Wm}^{-1}\text{K}^{-1})(2.0\text{m}^2)(20°\text{C} - 0°\text{C})}{8.0 \times 10^{-3}\text{m}} = 5000\text{W} = 5.0\text{kW}$$

The rate of heat loss through the window is 5.0kW.

Making Sense of the Answer: *Five kilowatts is the power of five small electric heaters running flat out, and it is leaking out of a single laboratory window. This is why glass windows are such a poor choice in cold climates and why double-glazing exists. The number is large because glass is a much worse insulator than its solid appearance suggests, and because 8mm is genuinely thin.*

Think Like a Physicist: *The first thing to do in any Fourier-law problem is to convert all lengths to metres. A thickness in millimetres dropped into the formula without conversion will give an answer that is wrong by a factor of a thousand, which is the kind of error that turns a 5kW leak into a 5W trickle, and a passing grade into a failing one.*

BINDER Example 2

A copper rod of length 2.0m and uniform cross-sectional area 4.0cm² is well lagged so that no heat escapes from its sides. One end is held at 100 °C and the other end is held at 20 °C. Take the thermal conductivity of copper as $380\text{Wm}^{-1}\text{K}^{-1}$. Find the rate of heat flow along the rod, and find the temperature gradient inside the rod.

Solution

Using Fourier's law for a uniform lagged rod:

$$\frac{dQ}{dt} = \frac{k\,A\,(\theta_1 - \theta_2)}{L}$$

Substituting:

$$\frac{dQ}{dt} = \frac{(380\text{Wm}^{-1}\text{K}^{-1})(4.0 \times 10^{-4}\text{m}^2)(100°\text{C} - 20°\text{C})}{2.0\text{m}} = 6.08\,\text{W}$$

The rate of heat flow along the rod is 6.08W.

For the temperature gradient: in a lagged uniform rod the temperature drops linearly from one end to the other. The gradient is therefore the same everywhere along the rod and equals the total temperature change divided by the length:

$$\frac{d\theta}{dx} = -\frac{(\theta_1 - \theta_2)}{L} = -\frac{(100 - 20)\text{K}}{2.0\text{m}} = -40\text{Km}^{-1}$$

The temperature gradient inside the rod is −40K per metre.

The minus sign confirms that temperature decreases as we move from the hot end (x = 0) toward the cold end (x = 2 m).

Making Sense of the Answer: *A 6W flow along a thin rod is small in absolute terms but large compared with what a wooden rod of identical dimensions would deliver. With k for wood near 0.15, the same rod made of wood would carry only about 0.0024 W under the same conditions, which is roughly 2500 times less.*

Think Like a Physicist: *In any uniform lagged rod, the temperature gradient is just the total temperature drop divided by the length. You almost never need calculus to find it. Save the dθ/dx machinery for situations where the geometry forces the gradient to vary, such as in radial heat flow through a sphere, which we will meet a few pages from now.*

REAL Example 3

A classroom in a Dar es Salaam secondary school has a single louvre-glass window of total area 1.5m² and thickness 4.0mm. On a hot afternoon, the outer surface of the glass reaches 38°C while the inner surface, cooled by a ceiling fan blowing air across it, sits at 30°C. Take k for glass as $1.0\text{Wm}^{-1}\text{K}^{-1}$.

(a) Calculate the rate at which heat is conducted into the classroom through the window.

(b) If the lesson lasts 80 minutes, calculate the total quantity of heat delivered into the classroom through that single window during the lesson.

Solution

(a) Heat flows from the hotter outside to the cooler inside. Applying Fourier's law:

$$\frac{dQ}{dt} = \frac{k\,A\,(\theta_1 - \theta_2)}{L} = \frac{(1.0\text{Wm}^{-1}\text{K}^{-1})(1.5\text{m}^2)(38°\text{C} - 30°\text{C})}{4.0 \times 10^{-3}\text{m}} = 3000\text{W} = 3.0\text{kW}$$

The rate of heat flow into the classroom is 3.0kW.

(b) Total heat delivered in 80 minutes.

Time t = 80 × 60s = 4800 s. Since the rate of heat flow is constant during the lesson:

$$Q = \frac{dQ}{dt} \times t = 3000\text{Js}^{-1} \times 4800\text{s} = 1.44 \times 10^7\,\text{J} = 14.4\text{MJ}$$

The total quantity of heat delivered through that single window is 1.44×10^7 J.

That 14.4MJ of heat delivered through one window in one lesson, equivalent to the energy released by burning nearly a kilogram of wood.

Making Sense of the Answer: *Three kilowatts coming through one window is roughly the heat output of three electric kettles boiling continuously. Multiply this by the four or five windows in a typical Tanzanian classroom and you have ten or fifteen kilowatts of solar heat being conducted in by the glass alone, which is exactly why classrooms in hot regions like Dar es Salaam grow steadily warmer through the morning and why architectural choices about glazing and shading actually matter for student comfort.*

Think Like a Physicist: *When a problem asks for total heat over a period of time, never forget to multiply the heat-flow rate by the time. Students often produce a beautifully correct rate in watts and then write that as the answer for the total energy. Watts are joules per second, not joules. Always finish the multiplication.*

With Fourier's law in our hands and the meaning of k properly understood, we are ready to look at what happens when the slab is not a single block of one material, when its sides are unable to keep heat in, or when its geometry is something other than a flat rectangle. Each of these complications introduces a new wrinkle, but the underlying law remains the same. Heat will keep flowing down the temperature gradient, at a rate determined by the conductivity, the geometry, and the temperature difference. The next subsection looks at the difference between a lagged bar and an unlagged bar, where the same rod can produce two completely different temperature profiles depending on whether or not its sides are insulated.

Lagged and Unlagged Bars

Fourier's law in the plain slab form assumes that all the heat that enters the hot face also leaves the cold face, with none escaping along the way. This is true for a slab whose sides are insulated, but it is dramatically false for a bare bar standing in open air. The same metal rod, with the same temperature applied at one end and the same temperature applied at the other, can produce two completely different temperature profiles depending on whether or not its sides are wrapped against the surrounding air. The two cases must be recognised on sight, because they obey different rules and demand different sketches.

The lagged bar

A bar is said to be **lagged** when *its sides are wrapped in a poor conductor (asbestos, wool, foam, or any insulating sleeve) so that no heat can escape sideways.* All the heat that enters the hot end is forced to travel along the bar and leave through the cold end. The lagging plays the same role as the four insulated walls of a pipe in fluid mechanics: ***it confines the flow to a single path.***

Because no heat is lost from the sides, the rate of heat flow at every cross-section along the bar is the same. If we pick any cross-section A along the length of the bar and ask how many joules per second are crossing it, the answer is identical to the answer at any other cross-section B. The bar acts as a one-dimensional conductor: heat enters at one end, travels straight through, and exits at the other end with no leakage.

This constancy of heat flow has a direct consequence. Applying Fourier's law between any two points on the bar:

$$\left(\frac{dQ}{dt}\right)_A = \left(\frac{dQ}{dt}\right)_B$$

Because the conductivity k is a property of the material, the area A is constant for a uniform bar, and the rate of heat flow is constant along the bar, the temperature gradient dθ/dx must also be constant everywhere inside the bar. A constant gradient means a straight-line temperature profile: temperature falls uniformly from θ_1 at the hot end to θ_2 at the cold end, with no curvature anywhere along the way.

The next figure shows both the lagged bar and the resulting temperature profile.

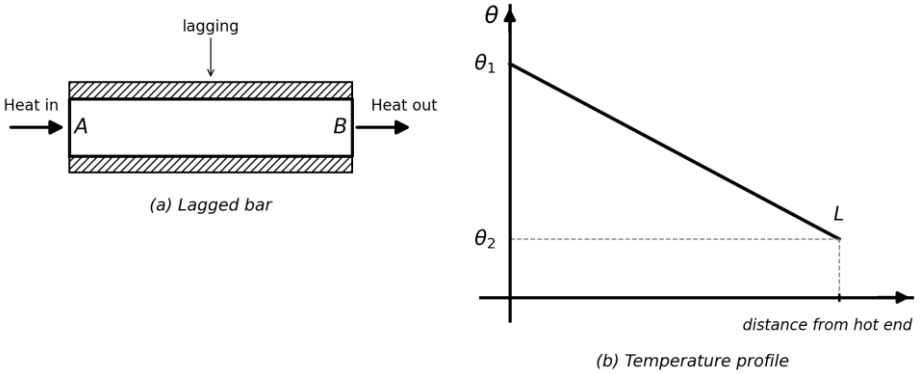


Figure: *(a) A lagged bar with insulation along its sides forces all heat to travel along the bar from end A to end B. (b) The temperature falls linearly from θ_1 to θ_2 along the length L of the bar, because the heat-flow rate (and hence the temperature gradient) is constant at every cross-section.*

Notice that the linear graph is what makes calculations on a lagged bar so straightforward. *Once the temperatures at the two ends are known, the temperature at any internal point can be read directly from the graph, or calculated from a single proportion:* a point one-third of the way from the hot end has a temperature one-third of the way from θ_1 to θ_2, and so on. No calculus is needed at any stage.

The unlagged bar

Now strip the insulation off the bar. The same metal, the same length, the same end temperatures, but this time the sides are bare and exposed to the surrounding air. Heat now has more than one road to travel. Some of it still flows along the bar from the hot end to the cold end, exactly as before, but some of it leaves through the sides at every cross-section, escaping into the air by convection and radiation. The bar is leaking heat continuously along its entire length.

This sideways leakage has two important consequences.

1) The heat-flow rate is no longer constant along the bar. At the hot end, where the bar is hottest, a large amount of heat enters; further along, after some heat has escaped from the sides, less heat is travelling through the bar; by the time the cold end is reached, only the small fraction that has not yet leaked out remains. The heat-flow rate decreases continuously from the hot end to the cold end.

2) The temperature gradient is no longer constant either. Because the heat-flow rate at any cross-section equals $-kA(d\theta/dx)$, and because the rate is now decreasing along the bar, the magnitude of the gradient must also be decreasing. Near the hot end, the bar carries a lot of heat and the gradient is steep; near the cold end, it carries very little and the gradient is gentle. The temperature profile is therefore not a straight line. It is a curve that falls steeply at the start and levels off toward the cold end. This shape is called **concave-up**: *it bows downward and flattens as the cold end is approached.*

The following figure shows the unlagged bar and its characteristic curved temperature profile, side by side with the lagged bar so the contrast can be seen at a glance.

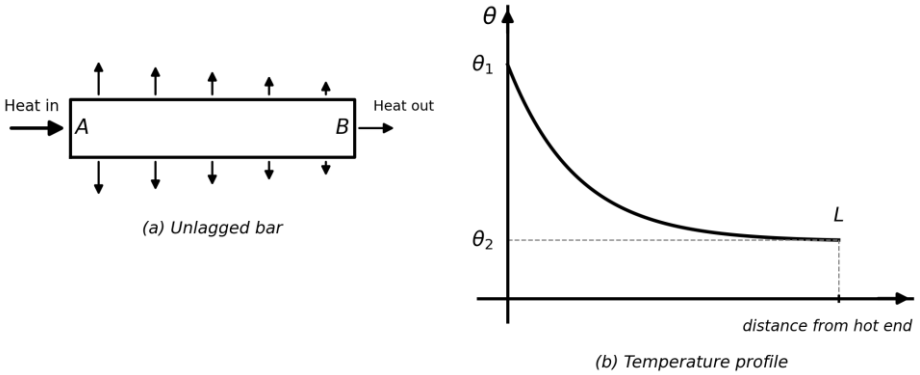


Figure: *(a) An unlagged bar loses heat from its sides along its entire length, so less heat reaches the cold end than enters the hot end. (b) The temperature profile is no longer linear; it falls steeply near the hot end and levels off toward the cold end.*

The contrast between the two bars is one of the most reliable indicators of physical understanding. A student who sketches a straight line for an unlagged bar has missed the point of the question, because the straight line is exactly the wrong answer. The straight line belongs to the lagged bar; the curve belongs to the unlagged one. Whenever a problem mentions that a bar has bare sides, or simply does not say it is lagged, the temperature profile must be drawn as a concave-up decay, never as a straight line.

Side-by-side comparison

The two cases are summarised in the table below for quick reference.

Feature	Lagged bar	Unlagged bar
Sides	wrapped in insulator	bare, exposed to air
Heat lost from sides	none	continuous along length
Rate of heat flow dQ/dt	constant at every cross-section	decreases from hot end to cold end
Temperature gradient dθ/dx	constant	magnitude decreases along bar
Temperature profile	linear (straight line)	concave-up (curved decay)

Time to put the two bars to work. Bring out your pen and let physics speak for itself.

BINDER Example 4

A copper bar of length 50cm and uniform cross-sectional area 4.0cm² is well lagged. The hot end is maintained at 100°C and the cold end at 30°C. Take the thermal conductivity of copper as $380\text{Wm}^{-1}\text{K}^{-1}$. Find the temperature at a point 20cm from the hot end.

Solution

Because the bar is lagged, the temperature profile is a straight line. The temperature at any internal point therefore lies on the straight line joining (0, 100°C) and (50cm, 30°C).

Given:
Length L = 50cm = 0.50m
Hot end $\theta_1 = 100°C$, cold end $\theta_2 = 30°C$
Distance from hot end x = 20cm = 0.20m
The total temperature drop is:

$$\theta_1 - \theta_2 = 100°C - 30°C = 70°C$$

Because the gradient is constant, the temperature drop over the first 20cm of the bar is in the same ratio as the distance covered:

$$\Delta\theta = \frac{x}{L} \times (\theta_1 - \theta_2) = \frac{0.20\text{m}}{0.50\text{m}} \times 70°C = 28°C$$

Subtracting this drop from the hot-end temperature gives the temperature at x = 0.20m:

$$\theta = \theta_1 - \Delta\theta = 100°C - 28°C = 72°C$$

The temperature at the point 20cm from the hot end is 72°C.

Making Sense of the Answer: *The point lies two-fifths of the way along the bar, so the temperature should drop by two-fifths of the total temperature drop. Two-fifths of 70°C is 28°C, leaving 72°C. The proportional answer matches the algebraic answer, which is the kind of double-check that catches arithmetic slips before they reach the marker.*

Think Like a Physicist: *In any lagged bar, the temperature at any internal point can be found by simple proportion, never by calculus. The straight-line profile is the gift that lagging gives you. Always check at the start of a heat-transfer problem whether the bar is lagged or not, because the answer changes completely between the two cases.*

HOT Example 5

A uniform unlagged metal bar has its hot end maintained at 100°C. At a steady state, the temperature 8.0cm from the hot end is 76°C, and the temperature 24cm from the hot end is 60°C.

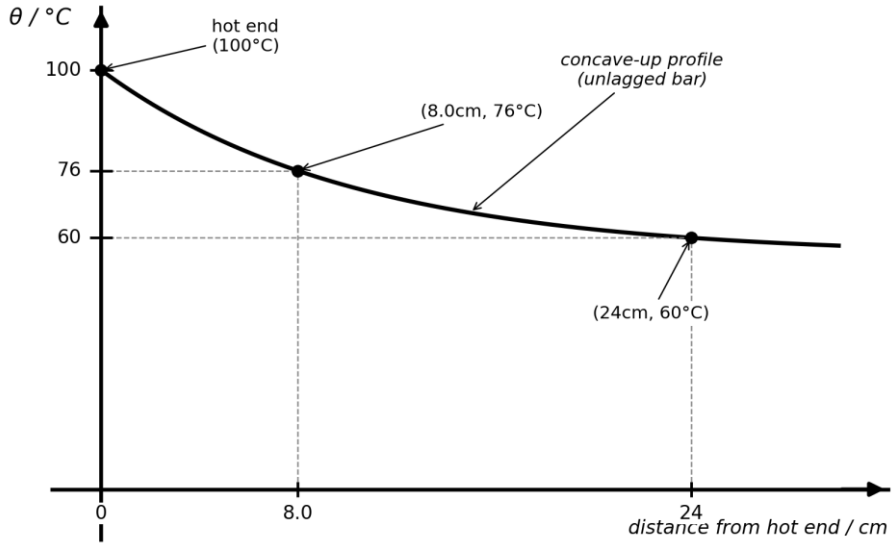
(a) Sketch the temperature profile of the bar.

(b) Calculate the average temperature gradient in the first region (0 to 8.0cm) and in the second region (8.0cm to 24cm), and compare them.

(c) Explain physically why the average gradient is steeper in the first region than in the second.

Solution

(a) The bar is unlagged, so the temperature profile is expected to be concave-up as sketched below.



(b) The average temperature gradient over a region is the temperature drop in that region divided by the length of the region.

First region (0 to 8.0cm):

$$\left(\frac{d\theta}{dx}\right)_1 = \frac{100°C - 76°C}{8.0cm} = \frac{24°C}{8.0cm} = 3.0°C/cm$$

Second region (8.0cm to 24cm), which is 16cm long:

$$\left(\frac{d\theta}{dx}\right)_2 = \frac{76°C - 60°C}{16cm} = \frac{16°C}{16cm} = 1.0°C/cm$$

The average gradient in the first region is 3.0°C/cm, three times steeper than the 1.0°C/cm in the second region.

(c) Near the hot end, very little heat has yet escaped from the sides, so almost all the heat that entered at the hot end is still travelling along the bar at the 8.0cm mark. A large rate of heat flow demands a large temperature gradient (because dQ/dt = −kA dθ/dx, and k and A are fixed for a uniform bar). Further along, the heat that was once travelling along the bar has been progressively leaking out of the sides, so by the time we reach the 8.0cm to 24cm region the bar is carrying much less heat. A smaller heat-flow rate requires a smaller temperature gradient, which is exactly what we calculated. The factor-of-three drop in average gradient between the two regions is direct evidence of the heat that has escaped sideways from the first 8.0cm of the bar.

Making Sense of the Answer: *Each region of the unlagged bar reports how much heat is still travelling along it at that point. A steep gradient near the hot end says "a lot of heat is still here". A gentle gradient further along says "most of the heat has already escaped sideways". The concave-up shape of the temperature profile is the bar's autobiography, written in the slope of its temperature curve.*

Think Like a Physicist: *When a problem asks about the gradient or the rate of heat flow at different positions along an unlagged bar, never assume they are equal. They are equal in a lagged bar; they are unequal in an unlagged one. The single word "unlagged" in a problem is a complete change of physical regime, and the entire approach to the problem must change with it.*

With the lagged and unlagged cases now firmly in hand, we are ready for the next complication. What if the bar is lagged but is not made of a single material? What if it is two different conductors joined end to end? This is the situation of a composite conductor, which the next subsection takes up.

Composite Conductors

So far we have applied Fourier's law to a single block of one material. Real walls, real roofs, and real cooking utensils are seldom this simple. The wall of a Tanzanian house may be built of brick on the outside, an air gap in the middle, and plaster on the inside; a cooking pot may have a thin steel base bonded to a thicker aluminium core; a domestic radiator is iron with a layer of paint on the outside. Each of these is a composite conductor, made of two or more materials joined together. The heat does not care that the materials are different. It simply flows from hot to cold, passing through whichever materials lie in its path. Our job is to predict the rate at which it does so.

Composite conductors come in two basic arrangements, and the analysis of each is genuinely beautiful because of how closely it parallels the analysis of resistors in an electric circuit. The two arrangements are slabs in **series** (one after the other, joined end to end so that *the heat must pass through them in turn*) and slabs in parallel (placed side by side so that *the heat can flow through either one independently*). Both arrangements are sketched in the following figure.

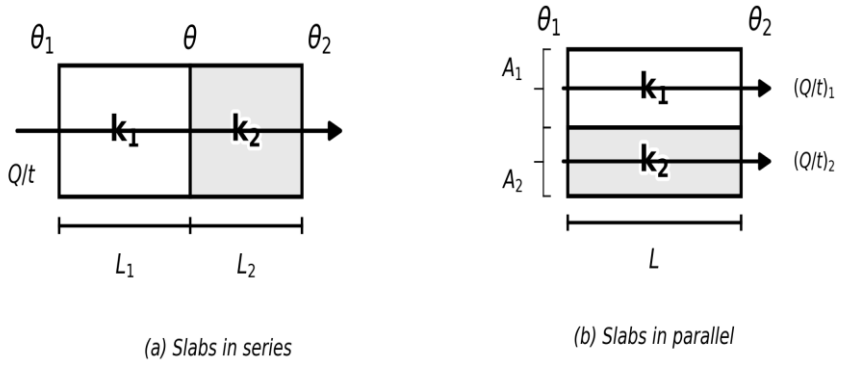


Figure: *(a) Two slabs of conductivities k_1 and k_2, lengths L_1 and L_2, joined end to end in series. The heat-flow rate Q/t is the same through both slabs. The interface temperature is denoted θ. (b) Two slabs of conductivities k_1 and k_2, cross-sectional areas A_1 and A_2, placed side by side in parallel. Each slab carries its own heat-flow rate, and the total heat flow is the sum of the two.*

Slabs in series

Consider two well-lagged slabs, both with the same uniform cross-sectional area A, joined end to end as shown in the figure. The first slab has thermal conductivity k_1 and length L_1; the second has thermal conductivity k_2 and length L_2. The outer face of the first slab is held at the higher temperature θ_1 and the outer face of the second slab is held at the lower temperature θ_2. We do not know the temperature at the interface where the two slabs meet, but we will call it θ and let the physics determine its value.

Because the slabs are well lagged, no heat escapes from the sides. All the heat that enters the first slab must therefore leave the second slab, and the rate of heat flow through the first slab must equal the rate of heat flow through the second slab. This is the key physical insight, and the entire derivation rests on it.

Let Q/t be the rate of heat flow through the composite. Applying Fourier's law to slab 1, where the temperature drops from θ_1 to θ across length L_1:

$$\frac{Q}{t} = \frac{k_1 A (\theta_1 - \theta)}{L_1}$$

Applying Fourier's law to slab 2, where the temperature drops from θ to θ_2 across length L_2:

$$\frac{Q}{t} = \frac{k_2 A (\theta - \theta_2)}{L_2}$$

These two equations express the heat-flow rate in two different ways, but the physics has insisted they must be equal. Setting them equal will eventually give us the interface temperature θ, but a more useful approach is to rearrange each equation so that the temperature drop sits alone, then add them.

Rearranging the first equation to isolate the temperature drop across slab 1:

$$\theta_1 - \theta = \frac{\left(\frac{Q}{t}\right) L_1}{k_1 A}$$

Rearranging the second equation to isolate the temperature drop across slab 2:

$$\theta - \theta_2 = \frac{\left(\frac{Q}{t}\right) L_2}{k_2 A}$$

Adding the two equations side by side, the unknown interface temperature θ cancels:

$$\theta_1 - \theta_2 = \frac{\left(\frac{Q}{t}\right)}{A} \times \left(\frac{L_1}{k_1} + \frac{L_2}{k_2}\right)$$

Solving for the heat-flow rate:

$$\frac{\mathbf{Q}}{\mathbf{t}} = \frac{\mathbf{A}(\boldsymbol{\theta_1} - \boldsymbol{\theta_2})}{\frac{\mathbf{L_1}}{\mathbf{k_1}} + \frac{\mathbf{L_2}}{\mathbf{k_2}}}$$

This is the working formula for two slabs in series. It generalises immediately: for any number n of slabs in series, the rate of heat flow through the composite is:

$$\frac{\mathbf{Q}}{\mathbf{t}} = \frac{\mathbf{A}(\boldsymbol{\theta_1} - \boldsymbol{\theta_2})}{\sum \frac{\mathbf{L_i}}{\mathbf{k_i}}}$$

where the sum runs over all slabs, and θ_1 and θ_2 are the outer-face temperatures of the entire composite (the hot and cold ends).

Recovering the interface temperature

Once the heat-flow rate Q/t is known, the interface temperature θ follows immediately from either of the original Fourier equations. Substituting Q/t back into the rearranged form for slab 1:

$$\boldsymbol{\theta = \theta_1 - \frac{\left(\frac{Q}{t}\right) L_1}{k_1 A}}$$

This will give a numerical value of θ that lies somewhere between θ_1 and θ_2. *The natural follow-up question is: closer to which temperature? Closer to the hot-end value θ_1, or closer to the cold-end value θ_2?*

A useful piece of physical insight: the interface temperature is closer to whichever side has the smaller temperature drop, and the smaller temperature drop occurs in whichever slab has the lower thermal resistance (high conductivity, short length). In short, the interface temperature is dragged toward the side that conducts heat more easily.

Thermal Resistance and the Electrical Analogy

The series formula we have just derived has a remarkable structure, and it deserves its own name and its own definition. Before we write them down, let us notice that this is not the first time we have met a situation where something flows, something drives the flow, and something opposes it. The same three-part story governs every electric circuit ever built, and physics rewards us for recognising the pattern.

In an electric circuit, a battery applies a potential difference V across a conductor. That potential difference drives a current I through the conductor. The conductor resists the flow, allowing only a limited current for each volt applied; the ratio V/I is called the resistance R, and the entire relationship is captured by Ohm's law V = IR. The units of resistance (volts per ampere, named the ohm) tell you the price in voltage you must pay to push one ampere through the device.

Now read that paragraph again, but replace every electrical word with its thermal counterpart. A **temperature difference $\theta_1 - \theta_2$** plays the role of **potential difference:** it is *the quantity that drives the flow.* The **rate of heat flow Q/t** plays the role of **current**: it is *the quantity that flows in response*, measured in joules per second, which is watts. And whatever plays the role of **resistance** should be the ratio of driving **temperature difference** to resulting **heat current**.

$$\textbf{Thermal resistance, } R = \frac{\textbf{Temperature difference}}{\textbf{heat current}} = \frac{\theta_1 - \theta_2}{\frac{Q}{t}}$$

So **thermal resistance** R of a slab is defined as *the ratio of the temperature difference across the slab to the rate of heat flow through it.*

Thermal resistance is therefore *a measure of opposition to heat current*, in exactly the sense that electrical resistance is a measure of opposition to electric current. A high-R slab demands a large temperature difference to produce a modest heat flow; a low-R slab lets heat flow freely for only a small temperature difference.

From Fourier's law:

$$\frac{Q}{t} = \frac{k A (\theta_1 - \theta_2)}{L}$$

Rearranging:

$$\frac{\theta_1 - \theta_2}{\frac{Q}{t}} = \frac{L}{k A} = R$$

So thermal resistance of a slab can also be defined as *the ratio of its thickness to the product of its thermal conductivity and cross-sectional area:*

$$\boldsymbol{R = \frac{L}{k A}}$$

The unit of R follows directly from the formula: metres divided by (W m^{-1} K^{-1} × m^2), which simplifies to **KW^{-1}.** A slab with high resistance impedes heat flow strongly; a slab with low resistance lets heat flow easily.

Every dependence in this formula matches what you would predict from the electrical analogy. A longer slab (larger L) has higher thermal resistance, just as a longer wire has more electrical resistance; heat must travel further and pushes against more molecules along the way. A broader slab (larger A) has lower thermal resistance, because heat has more parallel channels to flow through, just as a thicker wire carries current more easily. A material with higher thermal conductivity k has lower thermal resistance, in the same way a material with higher electrical conductivity has lower resistivity. Every intuition you already have about resistors in wires transfers, symbol by symbol, to slabs in heat flow.

The correspondence between heat and electricity is summarised in the table below.

Electrical quantity	Thermal counterpart
Potential difference V (volt)	Temperature difference $\theta_1 - \theta_2$ (kelvin)
Current I (ampere, = Cs^{-1})	Heat current Q/t (watt, = Js^{-1})
Resistance R = V/I (ohm, = VA^{-1})	Thermal resistance R = $(\theta_1 - \theta_2)/(Q/t)$ (KW^{-1})
Ohm's law: V = IR	Thermal Ohm's law: $(\theta_1 - \theta_2) = (Q/t)R$

The two columns are so perfectly parallel that every technique you have already learned for electrical circuits, including series addition of resistances, parallel reciprocals, voltage-divider reasoning, and Kirchhoff's laws, works without modification in heat-flow problems. The thermal version of a problem is simply the electrical version with relabelled symbols.

The combination rules inherit from Ohm's law

Because the mathematics of the thermal and electrical cases are identical, the combination rules are identical too. The series formula we derived earlier now reads, in resistance form:

$$R_{total} = R_1 + R_2 + \ldots + R_n$$

and the parallel rule, which we shall derive from Fourier's law in the next subsection, will have the familiar electrical reciprocal form:

$$\frac{1}{R_{total}} = \frac{1}{R_1} + \frac{1}{R_2} + \ldots + \frac{1}{R_n}$$

In either arrangement, the rate of heat flow through the composite is then given by thermal Ohm's law:

$$\frac{\mathbf{Q}}{\mathbf{t}} = \frac{\boldsymbol{\theta_1 - \theta_2}}{\mathbf{R_{total}}}$$

This is the same equation you would write for the current in an electrical circuit, with V replaced by $\theta_1 - \theta_2$ and I replaced by Q/t. Anyone who has solved a resistor-network problem has already solved every heat-flow problem they are likely to meet.

Thermal Insulance: the R-value of a Material

Thermal resistance, as we have defined it, depends on three things: the material (through k), the thickness (through L), and the cross-sectional area (through A) of the specific slab in question. The area dependence is convenient when we are working with a specific object such as the brass rod in the experiment or the copper base of a particular cooking pot, but it becomes inconvenient the moment we want to compare materials at a more general level.

Consider an insulation manufacturer preparing a product catalogue for fibreglass sheets. If they quoted a thermal resistance R, the number would only apply to pieces of a specific area; a customer cutting a larger or smaller sheet would get a different R. What the manufacturer actually wants to publish is a single number that characterises the insulating ability of the material-and-thickness combination by itself, independent of how big a piece the customer happens to order. That number is called **thermal insulance.**

The **thermal insulance r** of a material of given thickness *is a measure of its intrinsic ability to resist the transfer of heat, independent of the cross-sectional area of the specific sample.* For pure conduction through a uniform slab, its mathematical definition is:

$$\mathbf{r} = \frac{\mathbf{L}}{\mathbf{k}}$$

Its unit is $m^2\ K\ W^{-1}$.

Insulance and thermal resistance are related by the cross-sectional area of the actual slab:

$$\mathbf{R} = \frac{\mathbf{r}}{\mathbf{A}}$$

So thermal insulance is a pure material-and-thickness property: it tells you how much resistance content is packed into one unit of cross-section. Divide by the actual cross-section of the piece you have to recover the total thermal resistance R. A high-insulance material is a strong insulator; a low-insulance material lets heat through easily. The same piece of fibreglass insulation cut to different sizes has different thermal resistance but the same thermal insulance, which is exactly the property the catalogue wants to describe.

Why insulance captures the full heat-transfer story

We derived r = L/k for pure conduction through a single solid slab, but the concept generalises in a powerful way that justifies the broader name. The insulating performance of a real building wall is built up from several contributions in series: the conduction insulance of each solid layer (L/k for the brick course, L/k for the plaster, L/k for the foam board, and so on), plus additional insulance terms contributed by the thin stagnant-air films that cling to the wall's two external surfaces. At those surfaces, heat crosses partly by convection (air currents carrying energy away) and partly by radiation (infrared waves emitted

by the warm surface). These surface insulance terms are tabulated in standard references, in the same units m^2 K W^{-1}, and added to the conduction layers in series.

This is why the "R-value" quoted in insulation catalogues and building codes can always be added layer by layer to predict the performance of a wall without ever having to know the wall's area. The catalogue R-value is a **total thermal insulance**: a quantitative measure of how effectively a material or assembly can resist the transfer of heat through all three mechanisms (conduction, convection, and radiation) combined. Thermal insulance is the single currency in which the insulating performance of any material, assembly, or system is quoted; thermal resistance is simply that insulance applied to a specific area.

Thermal resistance and thermal insulance at a glance

Quantity	Formula	Units	What it tells you
Thermal resistance R	L/(kA)	K W^{-1}	Resistance of a specific slab of area A
Thermal insulance r	L/k	m^2 K W^{-1}	Material-and-thickness property, independent of area

The relationship between the two quantities is R = r/A, and the Fourier's-law formula for the heat-flow rate can be written in either currency. Using thermal resistance:

$$\frac{\mathbf{Q}}{\mathbf{t}} = \frac{\boldsymbol{\theta_1 - \theta_2}}{\mathbf{R}}$$

or, equivalently, using thermal insulance:

$$\frac{\mathbf{Q}}{\mathbf{t}} = \frac{\mathbf{A}\,(\boldsymbol{\theta_1 - \theta_2})}{\mathbf{r}}$$

Both forms are correct and give the same number. Use R when the specific slab dimensions are fixed up front and you want the Ohm's-law form; use r when the material property comes from a catalogue and you need to multiply by the wall area to get a heat flow.

Slabs in Parallel

Now consider two slabs of the same length L but possibly different cross-sectional areas A_1 and A_2, placed side by side as shown in figure drawn earlier (page 137). Both slabs span between the same two outer surfaces, which are held at temperatures θ_1 and θ_2. The heat now has two parallel paths to choose from: it can flow through slab 1 with conductivity k_1, or through slab 2 with conductivity k_2. In practice some heat takes each path, and the total heat-flow rate through the composite is the sum of the heat-flow rates through the two individual slabs.

Applying Fourier's law to each slab separately, with the same temperature difference $(\theta_1 - \theta_2)$ acting across both:

$$\left(\frac{Q}{t}\right)_1 = \frac{k_1 A_1(\theta_1 - \theta_2)}{L}$$

$$\left(\frac{Q}{t}\right)_2 = \frac{k_2 A_2(\theta_1 - \theta_2)}{L}$$

The total heat-flow rate is the sum:

$$\left(\frac{\mathbf{Q}}{\mathbf{t}}\right)_{\mathbf{total}} = \left(\frac{\mathbf{Q}}{\mathbf{t}}\right)_{\mathbf{1}} + \left(\frac{\mathbf{Q}}{\mathbf{t}}\right)_{\mathbf{2}} = \frac{(\mathbf{k_1\,A_1 + k_2\,A_2})\,(\boldsymbol{\theta_1 - \theta_2})}{\mathbf{L}}$$

In terms of thermal resistance, the parallel rule mirrors the electrical case exactly. The reciprocals of the individual resistances add to give the reciprocal of the total resistance:

$$\frac{1}{R_{total}} = \frac{1}{R_1} + \frac{1}{R_2} + \ldots + \frac{1}{R_n}$$

The two arrangements obey the same rules as their electrical counterparts: resistances in series add directly, and reciprocals of resistances in parallel add. Anyone who has wired a circuit can analyse a composite wall.

Summary of the Two Arrangements

The two arrangements and their working formulas are summarised below for quick reference.

Feature	Series	Parallel
Common quantity	heat-flow rate Q/t	temperature difference $(\theta_1 - \theta_2)$
What the slabs share	the same heat passes through each in turn	the two faces are at the same pair of temperatures
Total resistance	$R_{tot} = R_1 + R_2 + \ldots + R_n$	$1/R_{tot} = 1/R_1 + 1/R_2 + \ldots + 1/R_n$
Effective conductivity (equal lengths and areas)	$\mathbf{k_e = 2k_1k_2 / (k_1 + k_2)}$	$\mathbf{k_e = (k_1 + k_2)/2}$
Electrical analogue	resistors in series	resistors in parallel

Three examples. Bring your pen, your arithmetic, and your common sense. The physics will do the rest.

BINDER Example 6

A 10cm brass bar ($k = 109 Wm^{-1}K^{-1}$) is joined end-to-end to a 10cm copper bar ($k = 380 Wm^{-1}K^{-1}$), forming a composite of cross-section $6.0 cm^2$. The brass end is held at 100°C and the copper end at 20°C, and the entire arrangement is well lagged. Find: (a) the rate of heat flow along the composite, (b) the temperature at the brass-copper interface.

Solution

Two slabs in series between two fixed temperatures. This is a textbook Ohm's-law problem in thermal clothing.

Given:

Brass: $k_1 = 109 Wm^{-1}K^{-1}$, $L_1 = 10cm = 0.10m$

Copper: $k_2 = 380 Wm^{-1}K^{-1}$, $L_2 = 10cm = 0.10m$

Cross-section $A = 6.0 cm^2 = 6.0 \times 10^{-4} m^2$

Hot end $\theta_1 = 100°C$, cold end $\theta_2 = 20°C$

(a) *Compute the thermal resistance of each slab using R = L/(kA).*

Brass slab:

$$R_1 = \frac{L_1}{k_1 A} = \frac{0.10m}{(109 Wm^{-1}K^{-1})(6.0 \times 10^{-4} m^2)} = 1.529 KW^{-1}$$

Copper slab:

$$R_2 = \frac{L_2}{k_2 A} = \frac{0.10m}{(380 Wm^{-1}K^{-1})(6.0 \times 10^{-4} m^2)} = 0.439 KW^{-1}$$

Combine using the series rule:

$$R_{total} = R_1 + R_2 = 1.529 KW^{-1} + 0.439 KW^{-1} = 1.968 KW^{-1}$$

Apply thermal Ohm's law:

$$\frac{Q}{t} = \frac{\theta_1 - \theta_2}{R_{total}} = \frac{80K}{1.968 KW^{-1}} = 40.7W$$

The rate of heat flow along the composite is 40.7W.

(b) Interface temperature. The same heat current Q/t = 40.7W flows through the brass slab (series means equal current). The temperature drop across just the brass slab is given by thermal Ohm's law applied to that slab alone:

$$\theta_1 - \theta = \frac{Q}{t} \times R_1 = (40.7W)(1.529 KW^{-1}) = 62.2K$$

So the interface temperature is:

$$\theta = \theta_1 - 62.2K = 100°C - 62.2°C = 37.8°C$$

The interface temperature is 37.8°C.

Making Sense of the Answer: *The interface temperature has settled much closer to the cold end (20°C) than to the hot end (100°C). This is exactly what we expected: the brass slab has 3.5 times the thermal resistance of the copper slab (1.529 versus 0.439 KW⁻¹), so in a voltage-divider sense the brass slab "drops" 62.2K out of the 80K total while the copper slab drops only 17.8K. The interface therefore lies close to the temperature of the low-resistance side, which is the cold end.*

Think Like a Physicist: *In any series composite, before doing any calculation, identify which slab has the higher thermal resistance. The interface temperature will lie close to the temperature of the lower-resistance side. This single rule of thumb catches arithmetic errors instantly: if your interface temperature lands on the wrong side of the midpoint, you have made a slip somewhere in your calculation.*

REAL Example 7

A classroom at Usagara Secondary School has a brick exterior wall of total area $15 m^2$, thickness 150mm ($k = 0.70 Wm^{-1}K^{-1}$). Set into this wall is a single-pane glass louvre window of area $1.8 m^2$, thickness 4.0mm ($k = 1.0 Wm^{-1}K^{-1}$). On a hot Tanga afternoon, the outer surfaces of both the wall and the glass reach 40°C, while the inside surfaces are kept at 30°C by a ceiling fan. (a) Calculate the rate of heat gain through the glass window. (b) Calculate the rate of heat gain through the brick wall (the area outside the window). (c) Calculate the total rate of heat gain through the exterior wall including the window. (d) The window is only 11% of the total exterior-wall area; what percentage of the total heat gain does it contribute? Comment on what this tells us about building design in hot climates.

Solution

The window and the brick wall share the same inside and outside temperatures, so they sit in parallel as two independent heat-flow paths. We compute the thermal resistance of each path, then apply thermal Ohm's law to each path, then add the resulting heat currents.

Given:

Outside temperature $\theta_1 = 40°C$, inside temperature $\theta_2 = 30°C$, so $\Delta\theta = 10K$

Window: $A_w = 1.8 m^2$, $L_w = 4.0mm = 4.0 \times 10^{-3} m$, $k_w = 1.0 Wm^{-1}K^{-1}$

Brick wall (excluding window): $A_b = 15 m^2$, $L_b = 150mm = 0.150m$, $k_b = 0.70 Wm^{-1}K^{-1}$

(a) Heat gain through the glass window. Thermal resistance of the window pane:

$$R_w = \frac{L_w}{k_w A_w} = \frac{4.0 \times 10^{-3}\,\text{m}}{(1.0\text{Wm}^{-1}\text{K}^{-1})(1.8\text{m}^2)} = 2.22 \times 10^{-3}\,\text{KW}^{-1}$$

Applying thermal Ohm's law to the window:

$$\left(\frac{Q}{t}\right)_w = \frac{\Delta\theta}{R_w} = \frac{10\text{K}}{2.22 \times 10^{-3}\,\text{KW}^{-1}} = 4500\text{W}$$

(b) Heat gain through the brick wall. Thermal resistance of the brick:

$$R_b = \frac{L_b}{k_b A_b} = \frac{0.150\text{m}}{(0.70\text{Wm}^{-1}\text{K}^{-1})(15\text{m}^2)} = 1.43 \times 10^{-2}\,\text{KW}^{-1}$$

Again, applying thermal Ohm's law to the brick wall:

$$\left(\frac{Q}{t}\right)_b = \frac{\Delta\theta}{R_b} = \frac{10\text{K}}{1.43 \times 10^{-2}\,\text{KW}^{-1}} = 700\text{W}$$

(c) Total heat gain. In parallel, the heat currents add:

$$\left(\frac{Q}{t}\right)_{total} = \left(\frac{Q}{t}\right)_w + \left(\frac{Q}{t}\right)_b = 4500\text{W} + 700\text{W} = 5200\text{W}$$

The total rate of heat gain into the classroom through this exterior wall is about 5.2kW.

(d) Fraction of the total heat gain contributed by the window:

$$\frac{4500\text{W}}{5200\text{W}} = 0.87 = 87\%$$

The glass window, despite occupying only 11% of the exterior-wall area, contributes 87% of the total heat gain. The brick wall, although 8.3 times larger in area than the window, contributes only 13% of the total heat load entering through the exterior. Windows dominate the thermal performance of the wall.

Making Sense of the Answer: *The thermal-resistance numbers tell the whole story in one line: $R_b / R_w = 1.43 \times 10^{-2} / 2.22 \times 10^{-3} = 6.4$. The brick wall has about 6.4 times the thermal resistance of the window, so under the same 10K temperature difference it carries only 1/6.4 as much heat. This factor of 6.4 is the ratio of heat currents, and it is what pushes 87% of the total heat through the window even though the window covers only 11% of the wall area. Two effects combine to make the window so much more conductive than the wall it sits in: glass has a higher thermal conductivity than brick (1.0 versus 0.70), and the glass pane is 37 times thinner than the brick wall (4mm versus 150mm). The thinness matters much more than the material property, which is why even a large window is a drastically bigger heat path than the entire wall around it.*

Think Like a Physicist: *This is why traditional architecture in hot climates uses small windows, deep-set openings, wide overhangs, and heavy shutters, and why African vernacular buildings made of mud brick and thatch have always been cooler inside than modern glass-front imports. The physics is on the side of tradition. When architectural styles designed for cold climates (where large windows gather welcome passive solar heat) are imported into tropical cities, those same windows become unwanted heat-gain channels that work against the building's comfort. Understanding thermal resistance turns what would otherwise be cultural preference into quantified engineering: once the R-value of each element is known, the heat-flow rates through a wall can be predicted, and the wall can be designed to match the climate it stands in.*

HOT Example 8

Two slabs of the same material ($k = 0.50\text{Wm}^{-1}\text{K}^{-1}$) are placed in parallel between two surfaces held at 80°C and 20°C. Each slab has cross-sectional area 0.20m². The first slab is 3.0cm thick; the second is 5.0cm thick. (a) Calculate the heat-flow rate through each slab. (b) State which slab carries more heat and explain why. (c) Find the total heat-flow rate through the parallel arrangement. (d) If the same two slabs were instead placed in series (with the same outer-face temperatures), find the new total heat-flow rate.

Solution

Given:

Conductivity (both slabs) $k = 0.50\text{Wm}^{-1}\text{K}^{-1}$

Cross-section (each slab) $A = 0.20\text{m}^2$

Thicknesses $L_1 = 3.0\text{cm} = 0.030\text{m}$, $L_2 = 5.0\text{cm} = 0.050\text{m}$

Hot face $\theta_1 = 80°\text{C}$, cold face $\theta_2 = 20°\text{C}$, so $\Delta\theta = 60\text{K}$

(a) Slab 1 (3.0cm thick):

$$R_1 = \frac{L_1}{k A} = \frac{0.030\text{m}}{(0.50\text{Wm}^{-1}\text{K}^{-1})(0.20\text{m}^2)} = 0.30\text{KW}^{-1}$$

Slab 2 (5.0cm thick):

$$R_2 = \frac{L_2}{k A} = \frac{0.050\text{m}}{(0.50\text{Wm}^{-1}\text{K}^{-1})(0.20\text{m}^2)} = 0.50\text{KW}^{-1}$$

In parallel, both slabs experience the same $\Delta\theta = 60\text{K}$:

$$\left(\frac{Q}{t}\right)_1 = \frac{\Delta\theta}{R_1} = \frac{60K}{0.30KW^{-1}} = \mathbf{200W} \text{ (for first slab)}$$

$$\left(\frac{Q}{t}\right)_2 = \frac{\Delta\theta}{R_2} = \frac{60K}{0.50KW^{-1}} = \mathbf{120W} \text{ (for second slab)}$$

(b) First slab carries more heat (200W versus 120W). The two slabs are made of the same material and have the same cross-section and the same temperature difference, so the only thing that differentiates them is thermal resistance. The thinner slab has a lower R, and by thermal Ohm's law a lower R carries a larger Q/t under the same Δθ. Heat, like current, prefers the path of least resistance.

(c) Total heat-flow rate through the parallel arrangement. The resistances combine by the reciprocal-sum rule:

$$\frac{1}{R_{total}} = \frac{1}{R_1} + \frac{1}{R_2} = \frac{1}{0.30KW^{-1}} + \frac{1}{0.50KW^{-1}} = 5.33W\,K^{-1}$$

$$R_{total} = \frac{1}{5.33W\,K^{-1}} = 0.1875KW^{-1}$$

Applying thermal Ohm's law to the composite:

$$\left(\frac{Q}{t}\right)_{total} = \frac{\Delta\theta}{R_{total}} = \frac{60K}{0.1875KW^{-1}} = 320W$$

This matches the sum 200W + 120W = 320W obtained from part (a), as it must.

(d) For series:

$$R_{total} = R_1 + R_2 = 0.30KW^{-1} + 0.50KW^{-1} = 0.80KW^{-1}$$

Applying thermal Ohm's law:

$$\left(\frac{Q}{t}\right)_{total} = \frac{\Delta\theta}{R_{total}} = \frac{60K}{0.80KW^{-1}} = 75W$$

The series arrangement carries 75W, compared with 320W in parallel. The parallel arrangement carries more than four times as much heat as the series arrangement of the same two slabs.

Making Sense of the Answer: *The ratio total thermal resistance in series to that in parallel = 0.80 / 0.1875 = 4.27 tells the whole story in a single number. For the same Δθ, the heat current is inversely proportional to resistance, so rearranging the same two slabs from parallel to series increases the total resistance by a factor of 4.27 and decreases the heat current by the same factor. In parallel, the heat has two roads to choose from, and the roads share the load; in series, the heat is forced down one single road, with the full temperature drop spread over a longer combined path.* ***The parallel arrangement is always the better conductor, and the series arrangement is always the better insulator.*** *This is why insulation works by being thick (effectively a series arrangement of insulating layers) and why heat sinks work by being broad and finned (effectively a parallel arrangement of conducting paths).*

Think Like a Physicist: *When facing a composite-conductor problem, the very first move is to identify the arrangement.* ***Are the slabs sharing the same temperature difference but separate areas?*** *That is parallel: add the reciprocals of the resistances.* ***Are the slabs sharing the same heat flow but stacked end-to-end?*** *That is series: add the resistances directly. Once the arrangement is identified, the whole problem becomes an Ohm's-law calculation with thermal labels.*

With composite conductors now in our toolkit, we are ready for one more geometry. So far every conductor has been flat, with parallel faces and constant cross-sectional area. But many real conductors are not flat: a hot-water pipe, a steam main, a refrigerator wall curved around a corner, a metal sphere holding hot tea. For these we need radial conduction, where the cross-sectional area changes with position. That is the subject of the next subsection.

Radial Conduction

All the geometries we have considered so far share one thing in common: the cross-sectional area available for heat flow is the same at every point along the conducting path. A slab has a constant rectangular cross-section. A lagged bar has a constant cross-section along its length. A composite wall has a constant cross-section at every depth through the wall. In each case, the area A sits comfortably in Fourier's law as a single constant that can be pulled outside any algebra we need to do.

Real heat-transfer situations do not always cooperate. A hot-water pipe carries water at 80°C and sits in a room at 20°C. Heat flows outward from the inner surface of the pipe wall to the outer surface, but the area available for that heat to cross is not constant: at the inner surface the area is small (the circumference is short), while at the outer surface the area is larger (the circumference is longer). Every molecular shell of pipe wall through which the heat must pass has a slightly bigger surface area than the shell just inside it. The same is true of a spherical metal flask of hot tea, a steam-main lagging jacket, a cooking pot whose walls curve around a corner, and an electrical cable sheathed in insulation.

In these situations, the area A in Fourier's law depends on the radius r at which we evaluate it. We can no longer treat A as a constant; we must treat it as a function of position. The result is a Fourier equation that is a differential equation, and solving it requires a small amount of calculus. The reward is a clean formula for the heat-flow rate through any curved conductor, and a natural generalisation of the thermal-resistance concept we have already built.

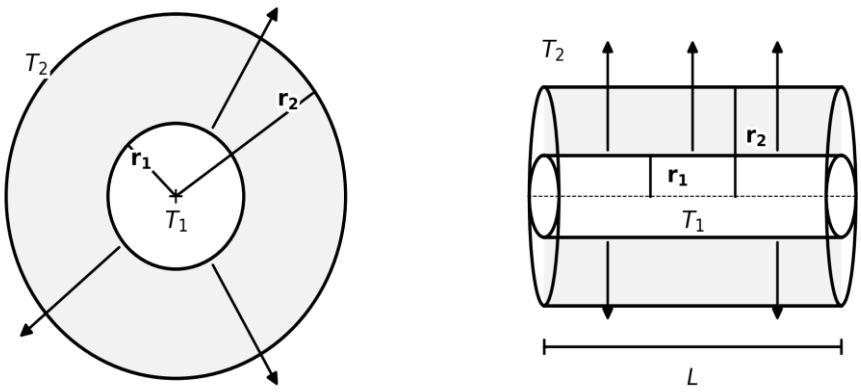


Figure: ***Two radial-conduction geometries.*** *(a) Concentric spheres: the hot inner sphere is at temperature T_1, the cold outer sphere is at T_2, and heat flows radially outward through every spherical shell of radius r between r_1 and r_2. (b) Coaxial cylinders of length L: the hot inner cylinder is at T_1, the cold outer cylinder is at T_2, and heat flows radially outward through every cylindrical shell of radius r between r_1 and r_2.*

Concentric Spheres

Consider two concentric spheres, as shown in the figure above. The inner sphere has radius r_1 and is held at the higher temperature T_1; the outer sphere has radius r_2 and is held at the lower temperature T_2. Between them is a solid shell of some material of thermal conductivity k (for instance, a spherical flask of tea surrounded by a layer of lagging). In a steady state, heat flows radially outward through every concentric shell of radius r lying between r_1 and r_2.

Consider one such shell, of inner radius r and thickness dr. This thin shell plays the role of a "slab" in Fourier's law, but with a few differences. Its area is the surface area of a sphere of radius r:

$$A(r) = 4\pi r^2$$

Its thickness is dr, and the temperature drop across it is dT (a small negative quantity, because temperature decreases as we move outward). Applying Fourier's law to this thin shell:

$$\frac{Q}{t} = -k\,(4\pi r^2)\,\frac{dT}{dr}$$

The minus sign confirms that heat flows in the direction of decreasing temperature. Because the arrangement is in a steady state, the heat-flow rate Q/t is the same through every concentric shell: all the heat that enters the inner surface of the lagging must emerge through its outer surface, with no accumulation anywhere in between. We can therefore treat Q/t as a constant during the integration and separate the variables:

$$\frac{Q}{t}\,\frac{dr}{r^2} = -4\pi k\,dT$$

Integrating from the inner surface ($r = r_1$, $T = T_1$) to the outer surface ($r = r_2$, $T = T_2$):

$$\frac{Q}{t}\int_{r_1}^{r_2}\frac{dr}{r^2} = -4\pi k\int_{T_1}^{T_2} dT$$

The integral on the left is $\int r^{-2}\,dr = -1/r$, and the integral on the right is straightforward:

$$\frac{Q}{t}\times\left(\frac{1}{r_1}-\frac{1}{r_2}\right) = 4\pi k\,(T_1 - T_2)$$

Solving for the heat-flow rate:

$$\boldsymbol{\frac{Q}{t} = \frac{4\pi k\,(T_1 - T_2)}{\frac{1}{r_1}-\frac{1}{r_2}}}$$

This is the working formula for steady-state radial conduction through a spherical shell.

Thermal resistance of a spherical shell

Writing the result in Ohm's-law form, $Q/t = (T_1 - T_2)/R$, the thermal resistance R of the spherical shell is read off as:

$$\boldsymbol{R_{sphere} = \frac{\frac{1}{r_1}-\frac{1}{r_2}}{4\pi k}}$$

Check the units: $1/r_1 - 1/r_2$ has units m^{-1}, $4\pi k$ has units $Wm^{-1}K^{-1}$, and their ratio comes out to KW^{-1} as it must. A spherical shell with small inner radius r_1 has large thermal resistance (if r_1 is small, $1/r_1$ is large), which is why *a tiny hot core loses heat*

slowly into a large outer volume: the heat is constrained to pass through a small inner surface before it can fan out into the expanding outer shells.

Coaxial Cylinders

Now consider two coaxial cylinders of length L, as shown in the figure on the previous page. The inner cylinder has radius r_1 and is held at the higher temperature T_1; the outer cylinder has radius r_2 and is held at T_2. Between them is a cylindrical shell of material of thermal conductivity k (for instance, the lagging around a hot-water pipe). In a steady state, heat flows radially outward through every cylindrical shell of radius r between r_1 and r_2.

Consider a thin cylindrical shell of inner radius r and thickness dr, sharing the same axis and the same length L as the outer cylinder. Its lateral surface area (the curved surface through which heat flows radially outward) is:

$$A(r) = 2\pi rL$$

Applying Fourier's law to this thin shell, with thickness dr and temperature drop dT across it:

$$\frac{Q}{t} = -k\,(2\pi rL)\,\frac{dT}{dr}$$

Separating variables, noting that Q/t is constant throughout the steady state:

$$\frac{Q}{t}\frac{dr}{r} = -2\pi kL\,dT$$

Integrating from $r = r_1$ (where $T = T_1$) to $r = r_2$ (where $T = T_2$):

$$\frac{Q}{t}\int_{r_1}^{r_2}\frac{dr}{r} = -2\pi kL\int_{T_1}^{T_2} dT$$

The integral on the left is $\int(1/r)\,dr = \ln(r)$, so:

$$\frac{Q}{t}\ln\left(\frac{r_2}{r_1}\right) = 2\pi kL\,(T_1 - T_2)$$

Solving for the heat-flow rate:

$$\boxed{\frac{Q}{t} = \frac{2\pi kL\,(T_1 - T_2)}{\ln\left(\frac{r_2}{r_1}\right)}}$$

This is the working formula for steady-state radial conduction through a cylindrical shell of length L. The natural logarithm in the denominator is the signature of cylindrical geometry: it arises because the area grows linearly with radius ($A \propto r$), so the reciprocal area (which appears in the integral) integrates to a logarithm rather than an algebraic expression.

Thermal resistance of a cylindrical shell

Writing the result in Ohm's-law form, $Q/t = (T_1 - T_2)/R$, the thermal resistance of the cylindrical shell is:

$$\boxed{R_{cyl} = \frac{\ln\left(\frac{r_2}{r_1}\right)}{2\pi kL}}$$

Notice that R scales inversely with the length L of the pipe, as we would expect: a longer pipe has more surface area, more cylindrical shells in parallel along its length, and therefore lower resistance to radial heat flow. The logarithm in the numerator means that doubling the outer radius does not double the resistance; it only adds a small amount ($\ln 2 \approx 0.69$) to whatever logarithm was already there. Thick lagging is subject to diminishing returns.

Summary of the Three Geometries

We have now derived Fourier's law for three different geometries: the uniform slab (from Section 3.1.1), the spherical shell, and the cylindrical shell. The three formulas are collected below for reference.

Geometry	Thermal resistance R	Heat-flow rate Q/t
Uniform slab (thickness L, area A)	$\frac{L}{kA}$	$\frac{kA(T_1 - T_2)}{L}$
Spherical shell (radii r_1, r_2)	$\frac{\frac{1}{r_1} - \frac{1}{r_2}}{4\pi k}$	$\frac{Q}{t} = \frac{4\pi k\,(T_1 - T_2)}{\frac{1}{r_1} - \frac{1}{r_2}}$
Cylindrical shell (length L, radii r_1, r_2)	$\frac{\ln\left(\frac{r_2}{r_1}\right)}{2\pi kL}$	$\frac{2\pi kL\,(T_1 - T_2)}{\ln\left(\frac{r_2}{r_1}\right)}$

The essential point is that in each case, the heat-flow rate can be written as (temperature difference) divided by (thermal resistance). The thermal resistance depends only on the geometry and the thermal conductivity of the material; it does not depend on the temperature difference. The Ohm's-law analogy survives the change of geometry unharmed.

Two curves, two calculations. Bring your calculator and pen.

BINDER Example 9

A hot-water pipe of inner radius 2.0cm and outer radius 2.5cm carries water at 80°C. The outer surface of the pipe is at 30°C. The pipe is made of copper ($k = 380 Wm^{-1}K^{-1}$) and has length 5.0m. Calculate the rate at which heat is conducted through the pipe wall.

Solution

This is the radial conduction through a cylindrical shell problem.

Thermal resistance of the cylindrical shell.

$$R_{cyl} = \frac{\ln\left(\frac{r_2}{r_1}\right)}{2\pi kL} = \frac{\ln(0.025/0.020)}{2\pi\,(380 Wm^{-1}K^{-1})(5.0m)} = 1.87 \times 10^{-5}\,KW^{-1}$$

Applying thermal Ohm's law:

$$\frac{Q}{t} = \frac{\Delta T}{R_{cyl}} = \frac{50K}{1.87 \times 10^{-5}\,KW^{-1}} = 2.68 \times 10^{6}\,W$$

The rate of heat conduction through the copper pipe wall is 2.68×10^6 W (about 2.7MW, or 2700kW).

Making Sense of the Answer: *A staggering answer: over two megawatts of heat flow through a 5-metre length of copper pipe. The number is large because copper is an excellent conductor (k = 380, one of the highest in the table) and the pipe wall is very thin (5mm of copper between hot water and cool outer surface). In practice a bare copper pipe would never maintain a 50K temperature difference across its wall for long: the outer surface would reach the water temperature within seconds, and the real heat loss would then be determined by convection from the outer surface to the room air, not by conduction through the copper wall. This is why hot-water pipes are always wrapped in lagging: the lagging is the actual bottleneck.*

Think Like a Physicist: *When computing radial conduction through a thin-walled metal conductor, expect a huge answer. The interesting physics is almost never the metal wall itself but what surrounds it. A good metal with a thin wall offers almost no resistance to heat flow, so its thermal resistance is tiny and the heat-flow rate it would carry under any realistic temperature difference is enormous. In any real installation, the temperature difference across the metal wall collapses to almost nothing within seconds, and the actual heat loss is then fixed by whatever sits outside the metal: lagging, paint, convection to the surrounding air. Whenever a problem gives you a metal-wall radial calculation and nothing else, the number you compute is a ceiling on the heat flow, not the heat flow a real pipe would actually deliver.*

HOT Example 10

An *alkasusu* seller at **Forodhani Gardens** in **Zanzibar** keeps his *alkasusu* hot in a spherical metal flask of inner radius 10cm and outer radius 11cm, made of steel ($k = 50 Wm^{-1}K^{-1}$). *Alkasusu* inside is at 85°C; the outer surface of the flask is at 25°C. (a) Calculate the rate of heat loss through the flask wall. (b) The seller wraps the flask in a 2.0cm-thick layer of felt ($k = 0.040 Wm^{-1}K^{-1}$), so the outer radius of the flask-plus-felt system becomes 13cm. Assuming the tea and the outer felt surface are still at 85°C and 25°C respectively, calculate the new rate of heat loss. (c) Comment on the ratio of the two heat-loss rates, and on which layer dominates the thermal resistance of the insulated system.

Solution

This is the radial conduction through concentric spheres problem. Part (a) uses a single spherical shell; part (b) uses two spherical shells in series.

(a) Rate of heat loss through the bare steel flask.

Thermal resistance of the steel shell:

$$R_{steel} = \frac{\frac{1}{r_1} - \frac{1}{r_2}}{4\pi k_{steel}} = \frac{\frac{1}{0.10m} - \frac{1}{0.11m}}{4\pi\,(50 Wm^{-1}K^{-1})} = \frac{0.909 m^{-1}}{628.3 Wm^{-1}K^{-1}} = 1.45 \times 10^{-3}\,KW^{-1}$$

Rate of heat loss through the bare flask:

$$\left(\frac{Q}{t}\right)_{bare} = \frac{\Delta T}{R_{steel}} = \frac{60K}{1.45 \times 10^{-3}\,KW^{-1}} = 41400W$$

The bare flask would lose heat at 41400W (about 41kW).

(b) Rate of heat loss through the flask + felt lagging. The two spherical shells are in series (the same heat crosses both in turn), so the total thermal resistance is the sum of the individual resistances.

Thermal resistance of the felt shell:

$$R_{felt} = \frac{\frac{1}{r_2} - \frac{1}{r_3}}{4\pi k_{felt}} = \frac{\frac{1}{0.11m} - \frac{1}{0.13m}}{4\pi\,(0.040 Wm^{-1}K^{-1})} = \frac{1.399 m^{-1}}{0.5027 Wm^{-1}K^{-1}} = 2.783 KW^{-1}$$

Total thermal resistance of the two shells in series:

$$R_{total} = R_{steel} + R_{felt} = 1.45 \times 10^{-3}\,KW^{-1} + 2.783KW^{-1} = 2.784KW^{-1}$$

Rate of heat loss through the lagged flask:

$$\left(\frac{Q}{t}\right)_{lagged} = \frac{\Delta T}{R_{total}} = \frac{60K}{2.784KW^{-1}} = 21.6W$$

The lagged flask loses heat at only 22W.

(c) Ratio of the two heat-loss rates:

$$\frac{\left(\frac{Q}{t}\right)_{bare}}{\left(\frac{Q}{t}\right)_{lagged}} = \frac{41400W}{21.6W} \approx 1900$$

The bare flask loses heat roughly 1900 times faster than the lagged flask.

Of the total thermal resistance of the lagged system, 2.783 out of 2.784 KW^{-1} (99.95%) comes from the felt alone. The steel flask wall, despite being made of a mediocre conductor (not as good as copper, but still a metal), contributes only 0.05% of the total resistance. Essentially all the insulation is provided by the felt.

Making Sense of the Answer: *The factor of 1900 is what makes tea flasks worth buying. Two centimetres of cheap felt converts a vessel that would cool in minutes into one that stays warm for hours. The physics behind this is the enormous conductivity ratio between steel and felt: k_{steel}/k_{felt} = 50 / 0.040 = 1250. Steel is 1250 times better at conducting heat than felt, and in series arrangement the resistance is dominated by whichever layer has the higher R. For two shells of comparable thickness, the one with the lower k wins the resistance battle by exactly the ratio of conductivities, which is why the felt controls everything even though it is only twice as thick as the steel.*

Think Like a Physicist: *This example illustrates a universal principle of heat-transfer engineering: whenever you have a layered system, the total resistance is dominated by the layer with the highest R. Adding more metal to the steel flask would barely help. Adding more felt to the lagging would help enormously. Real-world insulation design is therefore the art of identifying the weakest link in the series (the one with lowest R) and improving that one, rather than spending effort on layers that already dominate. The physics decides where your money should go.*

With radial conduction now added to our toolkit, we have completed the study of heat conduction in its three main geometries: slab, cylinder, and sphere. But conduction is only one of the three heat-transfer mechanisms introduced at the start of the chapter. The next section takes up the second: convection, in which a moving fluid carries heat bodily from place to place, and in which the mechanism of transport is no longer the passing of molecular vibrations from atom to atom but the bulk displacement of a heated material itself.

THERMAL CONVECTION

When Kipanga grabbed the steel spoon in the kitchen and felt the handle grow hot, he was feeling conduction: heat passing atom by atom through a solid that did not itself move. Now leave the kitchen for a moment and look up. The warmth you feel rising from a *sufuria* boiling on a stove, the breeze that pushes in from the ocean during a hot afternoon, the plume of warm air above a parked minibus in midday sun, even the rolling motion of water as it heats inside a kettle: all of these are examples of a different heat-transfer mechanism, one in which the hot material itself physically carries the energy from one place to another. We call it **convection.** By definition:

Thermal convection *is the process by which heat is transferred through a fluid (liquid or gas) by the actual movement of the fluid material itself from a region of higher temperature to a region of lower temperature.*

Convection is fundamentally different from conduction in one important way. In **conduction,** *the material stays where it is and energy migrates through it as a wave of molecular jostling.* In **convection,** *the material itself goes on a journey, and the energy rides along like a passenger.* You could blindfold a cup of water, put a lump of ice in one corner and a drop of heat from a match in the other, and track the water molecules as they moved: in conduction, almost none of them would move far from where they started. In convection, you would see them flowing around the cup in bulk, mixing hotter fluid with cooler fluid until the whole body came to a uniform temperature.

This distinction has an immediate consequence: *convection can only happen in materials whose molecules are free to move in bulk*. **Liquids** and **gases** qualify; solids do not. A metal bar can conduct heat but **cannot** convect it, *because its atoms are locked to their lattice sites*. A pot of water can both conduct and convect, but convection dominates quickly once the water starts moving. The air in a room above a radiator can barely conduct heat at all (air is one of the worst conductors in the k-table provided earlier), yet the room still warms throughout, because convection takes over the job.

Two Types of Convection

Convection currents can be set going by either of two different driving mechanisms, and it is important to keep them straight because the laws that govern the heat flow differ between them.

1) Natural convection

In **natural convection**, *the fluid moves because of density differences caused by uneven heating*. A pocket of fluid that has been warmed expands, becomes less dense than its surroundings, and is pushed upward by the buoyancy force: the same Archimedes effect that makes a cork float on water. At the top it eventually cools, becomes denser again, and sinks back down, displacing other warmed fluid as it falls. The result is a closed convection current, driven entirely by gravity acting on density gradients that the heating itself has created.

Every domestic example of natural heating works this way. A *sufuria* of water on a charcoal *jiko* shows convection currents in miniature, with hot water rising from the hot bottom of the pan and cooler water falling to take its place from the top. A room heated by a radiator or a fireplace does the same thing on a larger scale, as shown in the figure below.

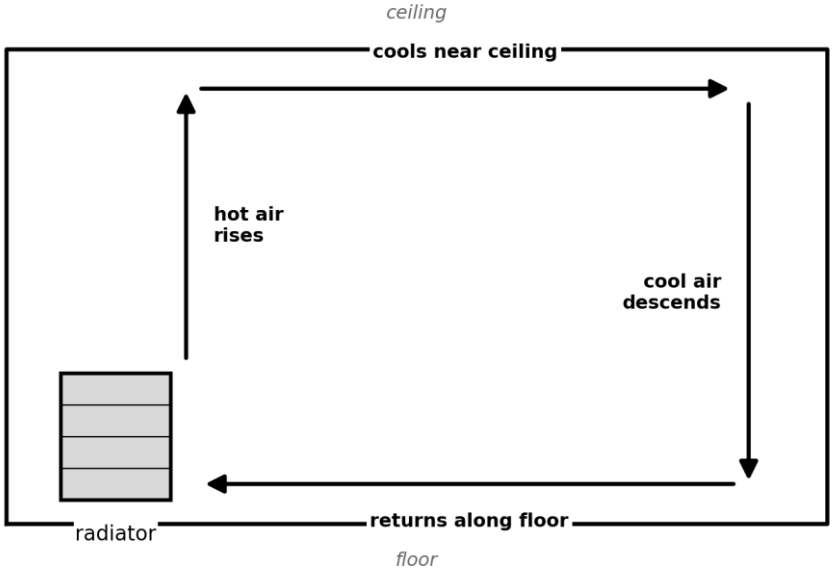


Figure: ***Natural convection in a room heated by a radiator****. Hot air rises from the radiator, travels across the ceiling (cooling as it goes by losing heat to the ceiling and walls), descends along the far wall, and returns along the floor to the radiator. The loop is driven entirely by the density difference between warm and cool air.*

On a much larger scale, the same principle drives the sea-land breeze cycle familiar to anyone who has lived along the Tanzanian coast. During a hot afternoon, the land heats faster than the ocean (because land has a much smaller specific heat capacity), warming the air above it. That warm air rises; cooler air over the ocean flows in at ground level to replace it; the result is the steady afternoon breeze from sea to land that keeps Dar es Salaam merely hot rather than intolerable. At night the effect reverses: the land cools faster than the ocean, the ocean now holds the warmer air, and the breeze blows from land out to sea. Trade winds, monsoons, and thunderstorm clouds are all driven by the same natural-convection mechanism operating on progressively larger scales.

2) Forced convection

In **forced convection**, *the fluid is driven by an external agent that imposes motion on it regardless of whether there is any temperature-induced density gradient.* A fan heater blows air past a hot element. A car radiator pumps coolant through metal channels past a hot engine block. The wind blows across a sweating face. A convection oven uses a built-in fan to circulate hot air past the food, cooking it more evenly than a still oven would. In every case an external force has taken over the job that gravity was doing in natural convection, and the fluid moves even where buoyancy would not carry it.

Forced convection is generally more effective at transferring heat than natural convection, because it is not limited by whatever modest density differences thermal expansion can create. A fan heater warms a room much faster than a radiator of the same power, because the fan forces air past the element at a speed far greater than natural buoyancy alone could achieve. This is why sweating cools the body more effectively in a breeze than in still air, and why hospitals use forced-air ventilation rather than trusting to natural circulation.

Factors Affecting the Rate of Convection

The rate at which a body loses heat by convection depends on three things. Any problem about convection can be approached by checking each of these three factors in turn.

Factor	What it means	Tanzanian example
Temperature difference between body and fluid	A larger excess temperature drives a steeper density gradient in natural convection and a larger heat flux in either mode.	A hot *sufuria* cools faster when the kitchen is cold than when it is already warm.
Properties of the fluid (viscosity, density, specific heat)	A low-viscosity, low-density, high-specific-heat fluid carries heat away more effectively per unit volume of flow.	A pot of tea cools faster in wind than in still water of the same temperature because air is less viscous and flows more freely.
Whether the convection is natural or forced	Forced convection is nearly always faster for a given temperature difference, because the fluid speed is not limited by buoyancy alone.	A person standing by an open minibus window on a moving *daladala* cools faster than a person sitting in a parked, windless minibus.

Dulong-Petit Law

When a hot body cools freely in still air (natural convection only) and the excess temperature above the surroundings is large, the rate of heat loss is found experimentally not to be proportional to the excess temperature itself but to a fractional power of it. The law was established by the French physicists Pierre Dulong and Alexis Petit in the early nineteenth century, and it states:

Under natural convection, the rate of heat loss from a hot body is directly proportional to the 5/4 power of the excess temperature of the body over its surroundings.

Mathematically:

$$\frac{dQ}{dt} \propto (\theta - \theta_s)^{5/4}$$

where θ is the temperature of the body and θ_s is the temperature of the surroundings. The fractional exponent 5/4 is a slightly steeper-than-linear dependence: doubling the excess temperature more than doubles the rate of heat loss (it multiplies it by $2^{5/4} \approx 2.38$).

The 5/4 power arises because natural convection depends on the strength of the convection current that the hot body creates, and that current is itself driven by the density difference between the warmed fluid near the body and the cool fluid far from it. A hotter body produces a larger density difference, hence a faster convection current, hence a more effective heat transfer. The precise exponent of 5/4 is an empirical result fitted to careful experiments; the important qualitative fact is that natural convection is a faster-than-linear process.

Dulong-Petit applies to the **natural-convection** regime with **large** excess temperatures (typically above about 50°C). It does **not** apply for **smaller excess temperatures**, and for the whole of the **forced-convection** regime at any excess temperature.

Newton's Law of Cooling

When a hot body cools either under **forced convection** (at **any** excess temperature) or under **natural convection** with a **small** excess (typically below about 30°C), the rate of heat loss is found to be directly proportional to the excess temperature itself. This was first stated by Isaac Newton in 1701 and has become one of the most-used laws in heat-transfer physics. The law states:

Under forced convection (at any excess temperature) or natural convection with a small excess temperature (below about 30°C), the rate of heat loss from a hot body is directly proportional to the excess temperature of the body over its surroundings.

Mathematically:

$$\frac{dQ}{dt} \propto (\theta - \theta_s)$$

The constant of proportionality depends on the body and its surroundings (its surface area, the fluid in contact with it, and whether the convection is natural or forced). Writing this constant as k, Newton's law becomes:

$$\frac{dQ}{dt} = k\,(\theta - \theta_s)$$

The full derivation of the corresponding temperature-versus-time expression, and the integration that predicts how long a hot cup of tea takes to cool, will be covered after the next section, which is devoted entirely to Newton's law and its consequences. For the purposes of this section it is enough to state the law, recognise its regime of validity, and understand why it sits alongside Dulong-Petit as one of two different laws for two different regimes of convection.

When does each law apply?

Knowing which of the two laws to use in a given situation is itself a piece of physics. The quick guide below captures the essentials.

Situation	Which law applies?
Forced convection, any excess temperature	Newton's law (linear in $\Delta\theta$)
Natural convection, small excess (roughly < 30°C)	Newton's law (linear in $\Delta\theta$)
Natural convection, large excess (roughly > 50°C)	Dulong-Petit law (proportional to $\Delta\theta^{5/4}$)
Natural convection, moderate excess (30°C to 50°C)	Transitional; neither law is exact

The thresholds given above are approximate and depend on the fluid involved, the geometry of the hot body, and the ambient conditions. In an A-level problem, assume Newton's law unless the excess temperature is very large and the convection is explicitly described as natural (no fan, no wind, no pump). If the problem mentions a fan or a moving fluid, assume Newton's law at any excess temperature.

Five examples. Three to warm up the concepts, two to drill the laws. Pens at the ready.

BINDER Example 11

A domestic iron radiator is installed on one wall of a room and heats the air around it. No fan is running. (a) Trace the complete convection-current loop that the radiator creates in the room, explaining what drives each segment of the loop. (b) Why does the upper part of the room warm up before the lower part? (c) How would the introduction of a small electric fan change the heating pattern?

Solution

(a) The radiator surface is hot and heats the air that touches it. The heated air expands, becomes less dense than the cooler air around it, and is pushed upward by buoyancy. This rising stream of hot air travels upward directly above the radiator until it reaches the ceiling. Along the way and at the ceiling, the air loses heat to the ceiling surface and the cooler surrounding air, becoming denser again. The now-cooler air moves along the ceiling away from the radiator, toward the far wall of the room. At the far wall, the air has cooled enough to become denser than the air below it; it descends down the far wall. Along the floor, this cool air is drawn back toward the radiator (partly by being displaced by more cool air coming down, partly by being pulled into the low-pressure region below the rising plume). When it reaches the radiator, it is heated again and the cycle repeats. The full loop is driven entirely by gravity acting on the density differences that the heating itself has created.

(b) The ceiling (upper part) warms before the floor (lower part) because the hot air rises first. Before the return flow along the floor has had time to bring cool air back to the radiator, heat has already collected at the ceiling. Also, the warm air is continuously replenished at the top of the room and continuously removed at the bottom, so the ceiling receives a steady delivery of fresh hot air while the floor receives a steady delivery of descending cool air.

(c) An electric fan converts the situation from natural convection to forced convection. The fan moves air at a speed much greater than buoyancy alone could produce, breaking up the orderly vertical temperature layering. Instead of the ceiling warming first and the floor last, the entire room is stirred into a roughly uniform temperature.

This is why ceiling fans are installed in Tanzanian homes: not to cool the air (they cannot change the air's temperature), but to mix it and to accelerate heat loss from the skin, which makes the body feel cooler.

Making Sense of the Answer: *The convection-loop trace is a chain of cause and effect, and each segment is driven by the result of the previous one. Hot air rises because it was heated. It cools because it travels along cooler surfaces. It descends because it became denser. It returns because it was drawn into the low-pressure zone left behind by the next rising plume. Every step is a consequence of the previous step. Any disruption to the chain (blocking the ceiling airflow with a curtain rail, obstructing the floor return with furniture) slows the whole loop.*

Think Like a Physicist: *When asked to trace a convection loop, do not simply say "hot air rises" and stop there. The full loop has four segments, and each segment has its own physical cause. A complete answer names all four:* ***(1) the rising plume*** *(driven by buoyancy),* ***(2) the sideways flow near the ceiling*** *(driven by displacement),* ***(3) the descending flow along the cool wall*** *(driven by increased density), and* ***(4) the return flow along the floor*** *(driven by continuity of flow). A one-segment answer is a one-segment understanding.*

REAL Example 12

During a hot afternoon in Mtwara, a breeze blows steadily from the ocean inland. At night the breeze reverses and blows from the land out to sea. (a) Explain both observations using the concepts of specific heat capacity and natural convection. (b) On a completely overcast day, when the sun is blocked and the temperatures of land and ocean stay roughly equal, there is essentially no breeze at all from either direction. Explain this observation.

Solution

(a) During the day, the sun heats both the land and the ocean. But land has a much smaller specific heat capacity than water (roughly $800 Jkg^{-1}K^{-1}$ for soil and rock versus $4200 Jkg^{-1}K^{-1}$ for water). Under the same solar input, the land temperature rises rapidly while the ocean temperature rises barely at all. By midday, the air over the hot land has been heated well above the temperature of the air over the cool ocean. The hot land-air is less dense; it rises. To replace it, cooler air flows in at ground level from the direction where the cool air is found, which is from the ocean. This is the sea breeze that makes Mtwara afternoons bearable.

At night, the sun no longer supplies heat, and both the land and the ocean cool by radiating heat into the sky. Again land has the smaller specific heat capacity, so it loses its daytime heat much faster than the ocean, which stays warm because it has so much stored thermal energy. By late night, the ocean is now warmer than the land. The air over the ocean is now the less-dense air; it rises. Cool air flows in at ground level from the now-cooler direction, which is from the land. This is the land breeze, blowing out to sea through the night and the early morning hours, which is why fishing dhows traditionally leave before dawn.

(b) On a completely overcast day, the sunlight is blocked by clouds and neither the land nor the ocean heats up significantly. With both surfaces at roughly the same temperature, the air above them is at roughly the same temperature everywhere, with no temperature difference to drive a convection current. Without a temperature difference there is no density gradient; without a density gradient there is no buoyant force to push air upward over one region; without such a force there is no flow in at ground level to replace the rising air. The atmosphere over land and ocean simply sits there motionlessly, and the breeze that would normally be driven by convection is absent. This is why overcast coastal days often feel strangely still and unpleasant, even at ocean-side locations that are normally windy.

Making Sense of the Answer: *The sea-land breeze system is a natural convection current on a continental scale, and every feature of it is explained by the same physics that governs a sufuria on a charcoal jiko. A temperature difference creates a density gradient; the density gradient drives a flow; the flow is completed into a closed loop by the continuity of the atmosphere. The system is larger than a kitchen but the physics is the same. This is also why "onshore flow" (sea breeze) is used by meteorologists as a reliable indicator that the land has been warming faster than the ocean that day, which itself is a standard marker of clear-sky insolation.*

Think Like a Physicist: *Whenever a real-world atmospheric phenomenon is described as a "wind" or "breeze", ask yourself what temperature difference is driving it. Almost every steady local wind on Earth (sea breezes, valley winds, monsoons, trade winds) has a temperature difference at its heart, and that temperature difference is almost always caused by differential heating of surfaces with different specific heat capacities. Identifying the temperature difference is the first step to understanding any wind system.*

BINDER Example 13

State the three factors that affect the rate of thermal convection from a hot body, and give a real-world Tanzanian example of each.

Solution

Factor 1: the temperature difference between the hot body and the surrounding fluid. A larger temperature difference drives a larger heat-flow rate by convection.

Example: a freshly boiled kettle (100°C) sitting in a 25°C kitchen loses heat much faster than the same kettle once it has cooled to 40°C. The first minutes of cooling feel urgently hot, while the last half hour feels barely warm; this is the temperature-difference factor at work.

Factor 2: the properties of the fluid in contact with the body, including its viscosity, density, and specific heat capacity.

Example: a hot *sufuria* placed outdoors on a windy day cools faster than the same *sufuria* placed in a bathtub of still water at the same air temperature, even though water has a much higher specific heat capacity than air. The reason is that air, with its low viscosity and low density, can be set moving far more easily than water; the convection current over the *sufuria* is much stronger in air than in still water, and the stronger current removes heat faster. Moving fluids beat still fluids, and the fluid's willingness to move is the determining factor.

Factor 3: whether the convection is natural or forced. Forced convection is generally much faster than natural convection.

Example: a bus conductor standing at an open window of a moving bus cools faster than a passenger sitting in the same bus with the windows closed. In the first case, the external motion of the bus creates a forced-convection flow of air across the conductor's skin, quickly sweeping away the warm boundary layer. In the second case, only natural convection is at work, and the boundary layer of warm air tends to cling to the skin, slowing heat loss.

Making Sense of the Answer: *The three factors are not independent of each other; in practice they combine. A large temperature difference in a moving fluid under forced convection produces the most rapid heat loss; a small temperature*

difference in a stagnant fluid under natural convection produces the slowest. Most everyday cooling situations involve some combination of all three, and estimating the cooling rate requires checking all three factors together.

Think Like a Physicist: *A useful way to remember the three factors is as a triangle: (1) the temperature difference is the driving force, (2) the fluid is the medium through which the heat must be carried, and (3) whether the flow is natural or forced is the delivery mechanism. Change any one of the three and the cooling rate changes; but always check all three before making a prediction. An answer that mentions only the temperature difference and ignores the fluid or the convection mode is an incomplete answer.*

BINDER Example 14

A hot plate at 200°C cools freely in still air at 25°C. Under these conditions the initial rate of heat loss is measured as 500W. Assuming that the cooling continues to obey the appropriate physical law throughout, estimate the rate of heat loss when the plate has cooled to 100°C.

Solution

The plate is cooling **freely** in **still air** (natural convection) and the initial excess temperature is $175\text{K}\ (200°\text{C} - 25°\text{C})$, well above $50°\text{C}$, so the **Dulong-Petit law** governs the cooling.

Interpreting the data:

Surrounding temperature $\theta_s = 25°\text{C}$

Initial plate temperature $\theta_1 = 200°\text{C}$, so initial excess $\Delta\theta_1 = 200°\text{C} - 25°\text{C} = 175°\text{C} = 175\text{K}$

Later plate temperature $\theta_2 = 100°\text{C}$, so later excess $\Delta\theta_2 = 100°\text{C} - 25°\text{C} = 75°\text{C} = 75\text{K}$

Initial rate $\left(\frac{\text{dQ}}{\text{dt}}\right)_1 = 500\text{W}$

Dulong-Petit law at two different excess temperatures:

$$\left(\frac{\text{dQ}}{\text{dt}}\right)_1 \propto (\Delta\theta_1)^{5/4}$$

$$\left(\frac{\text{dQ}}{\text{dt}}\right)_2 \propto (\Delta\theta_2)^{5/4}$$

Dividing one by the other cancels the proportionality constant:

$$\frac{\left(\frac{\text{dQ}}{\text{dt}}\right)_2}{\left(\frac{\text{dQ}}{\text{dt}}\right)_1} = \left(\frac{\Delta\theta_2}{\Delta\theta_1}\right)^{5/4}$$

Substituting the numerical values:

$$\frac{\left(\frac{\text{dQ}}{\text{dt}}\right)_2}{500\text{W}} = \left(\frac{75\text{K}}{175\text{K}}\right)^{5/4} = (0.4286)^{5/4}$$

$$\left(\frac{\text{dQ}}{\text{dt}}\right)_2 = 500\text{W} \times (0.4286)^{5/4} = 173\text{W}$$

The rate of heat loss when the plate has cooled to 100°C is about 173W.

Making Sense of the Answer: *The excess temperature has fallen by a factor of 175/75 = 2.33. Under a simple linear (Newton) law, the rate of heat loss would also fall by that same factor, to 500/2.33 = 215W. Under Dulong-Petit, with the steeper 5/4 exponent, the rate falls by a slightly larger factor of* $2.33^{5/4} = 2.89$*, to 173W. The difference of 215 vs 173W is a factor of about 1.24, which is the ratio of* $2.33^{1/4} \approx 1.24$*. This factor is small, but it is systematic: Dulong-Petit always predicts a somewhat faster cooling than Newton does, because the 5/4 exponent amplifies every excess temperature non-linearly.*

Think Like a Physicist: *When two physical laws both describe the same phenomenon in overlapping regimes, it is worth computing the same problem under both laws and comparing the answers. The ratio of the two predictions tells you how much the choice of law matters in the specific situation at hand. For large excess temperatures the two laws diverge measurably; for small excess temperatures they converge, which is exactly why Newton's law works for small excesses even though the true underlying physics is closer to Dulong-Petit.*

BINDER Example 15

Three hot bodies are cooling in a laboratory at 20°C. Body A is a cup of tea at 90°C sitting still on a table. Body B is the same cup of tea placed directly in front of a running desk fan. Body C is a bowl of soup at 35°C sitting still on the same table. For each body, decide whether its cooling is governed primarily by Newton's law of cooling or by the Dulong-Petit law, and justify your answer. No numerical calculation is needed.

Solution

Regime selection rests on two questions: is the convection natural or forced, and is the excess temperature large or small?

Body A: cup of tea at 90°C, no fan, excess temperature = 90 − 20 = 70K. This is natural convection (no external fluid driver) with a large excess temperature (well above the 50K threshold beyond which Dulong-Petit is the better description). Therefore, Dulong-Petit applies.

Body B: same cup of tea at 90°C, but with a fan blowing across it. The fan is an external agent driving the air past the tea, so the convection is forced regardless of the temperature difference. Under forced convection, Newton's law applies at any excess temperature. Therefore, Newton applies.

Body C: bowl of soup at 35°C, no fan, excess temperature = 35 − 20 = 15K. This is natural convection with a small excess temperature (well below the 30K threshold above which Dulong-Petit starts to take over). Therefore, Newton applies.

Making Sense of the Answer: *Of the three bodies, two obey Newton and only one obeys Dulong-Petit, which is a fair reflection of everyday life: most cooling situations involve either some forced airflow (a fan, a breeze, a breath) or a small excess temperature, and in both cases Newton's simpler linear law is the right tool. Dulong-Petit shines specifically for hot objects cooling in still air, which is less common in practice than it might seem. A laboratory setup with a hot plate on a bench and no air movement is one of the few situations where the 5/4 power genuinely describes the cooling better than a linear law would.*

Think Like a Physicist: *The regime-selection question is often the most important physical decision in a convection problem, and it should be made before any calculation begins. A student who plugs numbers into the wrong law gets a wrong answer no matter how carefully the arithmetic is done. Checking the two questions (forced or natural? large or small excess?) takes only a moment, and it settles which equation to use for the rest of the problem. Asking the right question first is half of solving it right.*

Thermal convection has now introduced the second of the three heat-transfer mechanisms, and has given us two laws (Dulong-Petit and Newton's) to describe the rate of heat loss by moving fluids. The detailed integration of Newton's law, which lets us predict how the temperature of a cooling body changes with time, is the subject of another section. Before that, however, we turn to the third mechanism: radiation, in which heat travels as electromagnetic waves and requires no medium at all. The next section takes up thermal radiation, from the glow of a charcoal *jiko* to the light of the Sun itself.

THERMAL RADIATION

Conduction transfers heat through a material by molecular jostling. Convection transfers heat by bodily moving the material itself. Both mechanisms require a material to be present. But the warmth of the Sun reaches us across 150 million kilometres of near-perfect vacuum, where there is no material at all. A third heat-transfer mechanism must exist, one that needs no medium and can carry energy through empty space. We call it **thermal radiation**.

Thermal radiation *is energy carried by electromagnetic waves, predominantly in the infrared part of the spectrum.* Every body whose temperature is above absolute zero emits thermal radiation continuously, regardless of its surroundings: a hot charcoal *jiko* radiates toward your face, but your face also radiates back toward the *jiko*. A warm Dar es Salaam pavement radiates toward the sky, and the sky radiates back toward the pavement. The only difference between a hot body and a cold one is how much radiation each emits per second. In this section we will learn exactly how much, and how the radiation depends on the body's temperature, its surface properties, and the temperature of its surroundings.

Blackbody Radiation and Emissivity

Some surfaces are better emitters and absorbers of radiation than others. A dull, dark, rough surface absorbs nearly all the radiation that falls on it and re-emits it efficiently; a shiny, polished, silvery surface reflects nearly all incoming radiation and emits very little of its own. To set a baseline for comparison, physicists define an idealised surface called a **blackbody**.

A **blackbody** *is a body that absorbs all radiation incident upon it, regardless of wavelength or angle of incidence.* Because it absorbs everything, it reflects nothing; and because it absorbs so efficiently, it also re-emits more efficiently than any other body at the same temperature. A **blackbody** is therefore both *a perfect absorber and a perfect emitter*.

No real surface is a perfect blackbody, but some come close. Soot, lampblack, and the interior of a furnace with a tiny observation hole are all excellent approximations. At the other extreme, a polished silver mirror is a very poor absorber and emitter.

To quantify how close a real surface comes to being a blackbody, we define a number called **emissivity**, given the symbol $\boldsymbol{\varepsilon}$ (the Greek letter epsilon). The **emissivity of a surface** *is the ratio of the power it actually radiates to the power a perfect blackbody of the same area and temperature would radiate.* **A blackbody has $\boldsymbol{\varepsilon = 1}$; a perfect reflector has $\boldsymbol{\varepsilon = 0}$.** Real surfaces fall between these two extremes.

Surface	Approximate emissivity ε
Perfect blackbody	1.00
Lampblack, soot	0.95
Dark rough paint	0.90
Human skin	0.80
Red brick, concrete	0.90
Oxidised copper	0.65
Polished copper	0.04
Polished aluminium	0.05
Polished silver	0.02

A useful rule of thumb emerges from this table: *good absorbers are good emitters, and good reflectors are poor emitters.* This is not a coincidence; it is a deep consequence of the microscopic physics of how electromagnetic waves interact with matter. *A surface that is efficient at absorbing a given wavelength of radiation must, by the same mechanism, be efficient at emitting that wavelength.* This principle, known as **Kirchhoff's law of thermal radiation**, is the reason why the inside of a thermos flask is silvered: the low emissivity of the silver surface reduces radiative heat loss, and the same low absorptivity reduces radiative heat gain. The silver acts as a radiation barrier in both directions.

Stefan-Boltzmann Law

The total power radiated (amount of radiation per unit time) by a body depends on its **temperature**, its **surface area**, and its **emissivity**. The quantitative relationship was discovered experimentally by Josef Stefan in 1879 and derived theoretically by Ludwig Boltzmann in 1884. The result is one of the most important equations in thermal physics.

Stefan-Boltzmann law states that:

The total power radiated by a body is proportional to the fourth power of its absolute temperature, to its surface area, and to its emissivity.

Mathematically:

$$\mathbf{P = \varepsilon \sigma A T^4}$$

where P is the radiated power in watts, ε is the emissivity (dimensionless, between 0 and 1), σ is the Stefan-Boltzmann constant $\sigma = 5.67 \times 10^{-8}\ \mathrm{Wm^{-2}K^{-4}}$, A is the surface area in m², and T is the absolute temperature in kelvin. The temperature must be in kelvin because the law involves an absolute power; the Celsius scale would give a meaningless result.

The fourth-power dependence is extremely steep. Doubling the temperature increases the radiated power by a factor of 16. Tripling the temperature increases it by a factor of 81. This is why a small increase in temperature produces a dramatically visible increase in brightness: a charcoal jiko at 800K glows dull red, but at 1600K (only twice the temperature) it glows blindingly white, radiating 16 times as much energy.

Net radiative power

A body does not radiate into nothing. It sits in an environment (surroundings, walls, air) that is itself at some temperature T_s and is itself radiating back toward the body. The net power lost by the body is the difference between what it emits and what it absorbs:

$$\mathbf{P_{net} = \varepsilon \sigma A\left(T^4 - T_s^4\right)}$$

If the body is hotter than its surroundings ($T > T_s$), the net power is positive and the body loses energy by radiation. If the body is cooler ($T < T_s$), the net power is negative and the body gains energy from its surroundings. At thermal equilibrium ($T = T_s$), the net power is zero: the body emits and absorbs at exactly the same rate.

Wien's Displacement Law

The Stefan-Boltzmann law tells us how much total power a body radiates, but it does not tell us how that radiation is distributed across different wavelengths. A glowing charcoal *jiko* looks reddish; the Sun looks yellowish-white; a welding arc looks bluish-white. As the temperature rises, the dominant colour shifts from red through yellow to blue-white. This shift is quantified by **Wien's displacement law.** It states that:

The wavelength at which a blackbody emits maximum spectral power is inversely proportional to its absolute temperature.

Mathematically:

$$\boldsymbol{\lambda_{max} = \frac{constant\ b}{T}} \text{ or } \boldsymbol{\lambda_{max}\, T = b}$$

where λ_{max} is the peak wavelength in metres, T is the absolute temperature in kelvin, and b is Wien's displacement constant, $b = 2.90 \times 10^{-3}$ m K. The product λ_{max} T is the same for every blackbody regardless of temperature: a hotter body shifts its peak to shorter wavelength so that the product stays constant.

The displacement law is beautifully illustrated by the spectral emission curves shown in the following figure.

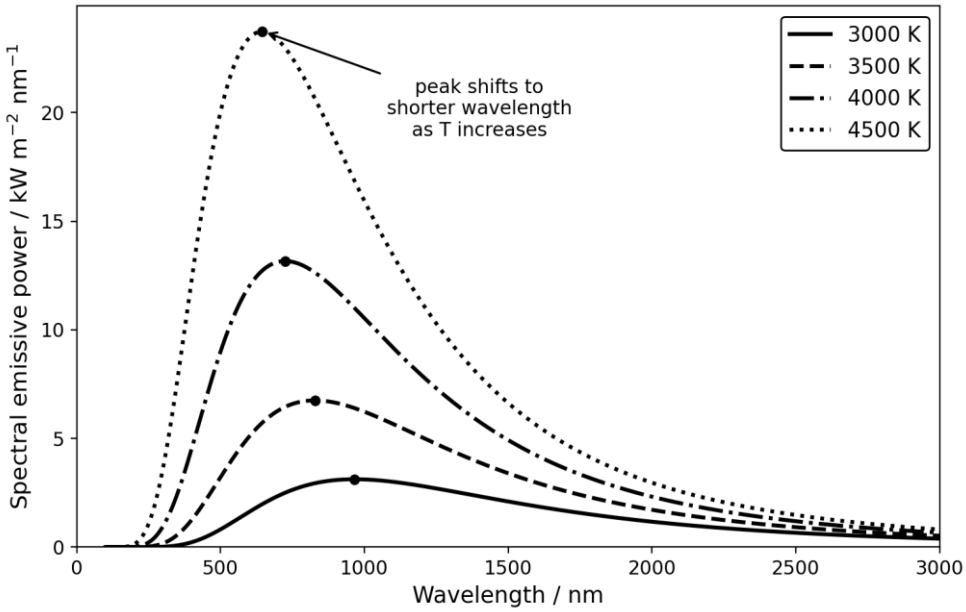


Figure: ***Spectral emission curves for blackbodies at four temperatures from 3000K to 4500K****. As the temperature increases, the peak shifts to shorter wavelength (Wien's law) and the total area under the curve grows rapidly (Stefan-Boltzmann's law). The dots mark the peak wavelength at each temperature.*

Key Observations from the Spectral Curves

Five features of the graph are worth studying carefully. Each one tells a different piece of the radiation story.

1. The peak shifts to shorter wavelength as temperature rises

At 3000K the peak sits deep in the infrared, beyond 900nm. At 4500K it has moved to about 640nm, near the red end of the visible spectrum. This leftward shift is Wien's displacement law in action: the product of peak wavelength and temperature stays constant, so a higher temperature forces the peak to a shorter wavelength. This is why a heating body changes colour from invisible (infrared) through dull red, orange, yellow, and eventually to white and blue-white as its temperature rises.

2. The peak grows taller at higher temperatures

The 4500K curve towers over the 3000K curve. This is not a separate effect from the wavelength shift; it is a physical consequence of it. All electromagnetic radiation travels at the speed of light, so *shorter wavelength means higher frequency. The energy carried by each photon is proportional to its frequency*. When the peak moves to shorter wavelength, it moves to a part of the spectrum where each photon is individually more energetic. A hotter body therefore radiates more power at its peak for two reasons working together: it emits more photons per second **and** each photon near the new peak carries more energy than the photons near the old peak. The taller peak is a direct consequence of the shift to shorter wavelength.

3. The total area under each curve grows rapidly with temperature

The area under the spectral curve represents the total power radiated per unit surface area. This area grows as the fourth power of the absolute temperature, which is Stefan-Boltzmann's law visible on the graph. Doubling the temperature increases the area by a factor of 16. The difference between the 3000K and 4500K curves is a ratio of $(4500/3000)^4 = (1.5)^4 = 5.06$, meaning the 4500K body radiates about five times as much total power per unit area as the 3000K body.

4. The peak height scales as the fifth power of temperature

This follows from observations 2 and 3 acting together. Stefan-Boltzmann says the total area grows as the fourth power of T. Wien says the peak shifts to shorter wavelength, which compresses the curve horizontally into a narrower band. More area squeezed into a narrower band forces the curve to grow taller. The combined effect is that the peak height is proportional to the fifth power of T: doubling the temperature makes the peak 32 times taller. This is why the height difference between curves on the graph is even more dramatic than the area difference.

5. The colour sequence of a heated body

As a body is heated continuously from room temperature, its dominant emission progresses through the following sequence: invisible (infrared only) at low temperatures, then dull red, bright red, orange, yellow, white (a mixture of all visible wavelengths), and finally blue-white at extremely high temperatures. The Sun's surface at about 5800K emits a spectrum that peaks in the visible range, which is why sunlight looks approximately white. A tungsten light-bulb filament at about 2800K peaks in the infrared, which is why incandescent bulbs are poor light sources: most of their radiated energy comes out as invisible heat rather than visible light.

The ideas have finished cooking; now let us sit down and enjoy them through three worked examples.

BINDER Example 16

A patient lies in an examination room at 16°C. The patient's skin temperature is 34°C, their body surface area is 1.6m², and the emissivity of human skin is 0.80. (a) Calculate the total power radiated by the patient. (b) Calculate the power the patient absorbs from the surroundings. (c) Find the net rate of radiative heat loss.

Solution

In this problem, Stefan-Boltzmann law applied twice: once for emission, once for absorption.

Given:

Skin temperature $T = 34°C = 307K$, surroundings $T_s = 16°C = 289K$

Surface area $A = 1.6m^2$, emissivity $\varepsilon = 0.80$

Stefan-Boltzmann constant $\sigma = 5.67 \times 10^{-8}\ Wm^{-2}K^{-4}$

(a) Total power radiated by the patient:

$$P_{emit} = \varepsilon\sigma AT^4 = (0.80)(5.67 \times 10^{-8}\ Wm^{-2}K^{-4})(1.6m^2)(307K)^4 = 645W$$

The total power radiated by the patient is $645W$.

(b) Power absorbed from the surroundings:

$$P_{abs} = \varepsilon\sigma AT_s^4 = 0.80 \times 5.67 \times 10^{-8}\ Wm^{-2}K^{-4} \times 1.6m^2 \times (289K)^4 = 506W$$

The power the patient absorbs from the surroundings is $506W$.

(c) Net rate of radiative heat loss:

$$P_{net} = P_{emit} - P_{abs} = 645W - 506W = 139W$$

The net rate of radiative heat loss is 139W.

The patient loses about 139W of heat by radiation alone.

Making Sense of the Answer: *139 watts is a significant fraction of the body's resting metabolic rate (roughly 80 to 100W for an average adult at rest). The fact that radiative heat loss alone exceeds the resting metabolic rate means the body must rely on other sources of internal heat (muscular activity, food digestion, shivering) to maintain its temperature in a 16°C room. This is why patients in cold examination rooms feel cold very quickly even without wind: radiation alone is pulling heat away faster than the body at rest can generate it.*

Think Like a Physicist: *In any Stefan-Boltzmann problem, it is better to compute both the emitted and absorbed powers separately before subtracting. It is tempting to jump directly to the net formula, but computing each term individually gives you a sanity check: the emitted power should always exceed the absorbed power for a body hotter than its surroundings, and the absorbed power should never be negative (it represents real energy arriving from the environment). A negative absorbed power means a sign error in your temperature conversion.*

BINDER Example 17

A furnace radiates maximum energy at a wavelength of 1200Å. A star emits its maximum radiation at 4800Å. (a) Calculate the surface temperature of each body. (b) Find the ratio of the power radiated per unit area by the furnace to that by the star.

Solution

Note: $1Å = 10^{-10}m$.

(a) Surface temperatures from Wien's displacement law:

For the furnace, $\lambda_{max} = 1200Å = 1200 \times 10^{-10}\ m = 1.20 \times 10^{-7}\ m$:

$$T_f = \frac{b}{\lambda_{max}} = \frac{2.90 \times 10^{-3}\ mK}{1.20 \times 10^{-7}\ m} = 24167K$$

For the star, $\lambda_{max} = 4800Å = 4.80 \times 10^{-7}\ m$:

$$T_s = \frac{2.90 \times 10^{-3}\ mK}{4.80 \times 10^{-7}\ m} = 6042K$$

The furnace is at $24167K$ and the star is at $6042K$ (comparable to the Sun's surface temperature of 5800K).

(b) From Stefan-Boltzmann law:

$$P = \varepsilon\sigma AT^4$$

The power per unit area is:

$$\frac{P}{A} = \varepsilon\sigma T^4$$

Both a furnace cavity and a stellar surface emit radiation that closely matches the Planck blackbody spectrum, so we take $\varepsilon = 1$ for both. The ratio is:

$$\frac{\left(\frac{P}{A}\right)_f}{\left(\frac{P}{A}\right)_s} = \left(\frac{T_f}{T_s}\right)^4 = \left(\frac{24167K}{6042K}\right)^4 = (4.0)^4 = 256$$

The furnace radiates 256 times as much power per unit area as the star.

Making Sense of the Answer: *The furnace is only 4 times hotter than the star, but it radiates 256 times as much power per unit area. The fourth-power law amplifies every temperature difference dramatically. This is why industrial furnaces and welding arcs consume such extraordinary amounts of energy: the power they radiate scales as the fourth power of their temperature, and at several thousand kelvin the losses become prodigious.*

Think Like a Physicist: *A useful way to estimate Stefan-Boltzmann ratios quickly: if the temperature ratio is n, the power ratio is n to the fourth. For n = 2, the ratio is 16; for n = 3, it is 81; for n = 4, it is 256; for n = 10, it is 10000. These round number let you predict the answer to any comparative radiation problem before touching a calculator.*

BINDER Example 18

A piece of metal loses $255Js^{-1}$ by radiation when its temperature is 1200K and the surroundings are at 300K. Find the rate of heat loss by radiation when the metal cools to 600K, with the same surroundings.

Solution

Interpreting the data:
At state 1: $T_1 = 1200K$, $T_s = 300K$, $P_1 = 255Js^{-1}$
At state 2: $T_2 = 600K$, $T_s = 300K$
Net Stefan-Boltzmann equation at each state:

$$P_1 = \varepsilon\sigma A\left({T_1}^4 - {T_s}^4\right)$$
$$P_2 = \varepsilon\sigma A\left({T_2}^4 - {T_s}^4\right)$$

Dividing cancels $\varepsilon\sigma A$:

$$\frac{P_2}{P_1} = \frac{{T_2}^4 - {T_s}^4}{{T_1}^4 - {T_s}^4}$$

Substituting (the K^4 units cancel in numerator and denominator):

$$\frac{P_2}{255Js^{-1}} = \frac{(600K)^4 - (300K)^4}{(1200K)^4 - (300K)^4}$$

Solving for P_2:

$$P_2 = 15.0Js^{-1}$$

The rate of heat loss at 600K is $15.0Js^{-1}$.

Making Sense of the Answer: *The metal has halved in temperature (1200K to 600K), but the net radiative loss has fallen from 255W to 15W, a factor of 17. This is far steeper than the factor of 16 you would expect from the fourth-power law applied to the body temperature alone, because the net formula also includes the surroundings term. At 1200K, the surroundings (300K) are cold compared to the body and barely matter. At 600K, the surroundings are only a factor of 2 away, and their contribution to the absorbed power is now significant: the body's emission has fallen dramatically (by a factor of 16), but its absorption from the surroundings has stayed the same (the surroundings are still at 300K). The net loss shrinks even faster than the gross emission.*

Think Like a Physicist: *In any comparative radiation problem, always compute the full net expression rather than approximating by dropping the surroundings term. The surroundings term may be negligible when the body is very hot, but it becomes dominant as the body approaches the surroundings temperature. Dropping it early is a trap that produces the right-looking answer (a clean power of 4 ratio) at high temperatures but the wrong answer at moderate temperatures.*

Stefan-Boltzmann and Wien together describe the total power and the spectral distribution of blackbody radiation. But radiation is a two-way street: every body emits and every body absorbs, simultaneously and continuously. The next subsection takes up this two-way picture, called Prevost's theory of heat exchange, and shows how it explains phenomena from premature-baby incubators to thermos flasks. After that we will derive the solar constant, predict Earth's equilibrium temperature, and show that Newton's law of cooling is a special limiting case of Stefan's law.

Prevost's Theory of Heat Exchange

It is tempting to think that a hot body radiates and a cold body absorbs. **Prevost's theory**, stated by Pierre Prevost in 1791, corrects this picture in an important way.

According to **Prevost's theory of heat exchange:** *Every body continuously emits thermal radiation and simultaneously absorbs thermal radiation from its surroundings, regardless of whether the body is hotter than, cooler than, or at the same temperature as its surroundings.*

In other words, the theory suggests that: *Emission and absorption are not alternatives; they happen together, at all times, for every body above absolute zero.* A body does not "stop emitting" when it reaches thermal equilibrium. It simply reaches a state in which its rate of emission equals its rate of absorption, so the net exchange is zero and its temperature stays constant.

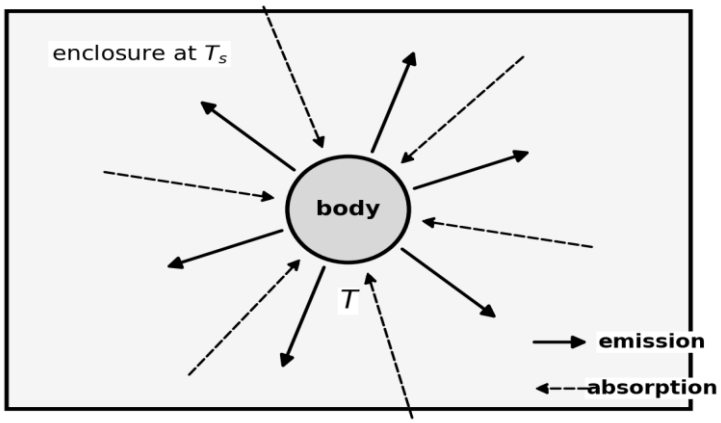


Figure: ***Prevost's theory of heat exchange.*** *A body at temperature T sits inside an enclosure at temperature T_s. Solid arrows represent emission (outward from the body); dashed arrows represent absorption (inward from the enclosure walls). Both processes occur simultaneously and continuously, regardless of which is hotter.*

Consequences of Prevost's theory

1. At thermal equilibrium ($T = T_s$), the body still radiates and still absorbs. The two rates are equal, so there is no net heat transfer and the temperature does not change. Equilibrium is a dynamic balance, not a cessation of activity.
2. A premature baby in a hospital incubator may be kept in warm air at, say, 35°C, yet still lose dangerous amounts of heat by radiation if the walls of the incubator are cold (for instance, if the incubator is near a cold window). The baby's skin radiates toward the cold walls, and the cold walls radiate back less than the baby emits. The net radiative loss can be large enough to cause hypothermia, even though the air temperature is perfectly adequate. Modern neonatal incubators solve this problem by using double-walled enclosures with warm inner walls, so that the wall radiation arriving at the baby's skin is close to what the baby emits.
3. A vacuum (thermos) flask reduces radiative heat transfer by silvering the inner and outer walls. The silver surface has very low emissivity ($\varepsilon \approx 0.02$), so the hot liquid inside emits very little radiation outward, and the cold environment outside sends very little radiation inward. Both directions are suppressed by the same low emissivity, exactly as Prevost's theory predicts: reducing emission necessarily reduces absorption by the same factor.

The Solar Constant and Earth's Equilibrium Temperature

The Sun is the most important radiating body in our daily lives. Everything from the temperature of the atmosphere to the direction of the sea breeze depends on how much solar radiation reaches Earth's surface. In this subsection we derive the solar constant from Stefan-Boltzmann's law and use it to predict what Earth's temperature would be without an atmosphere.

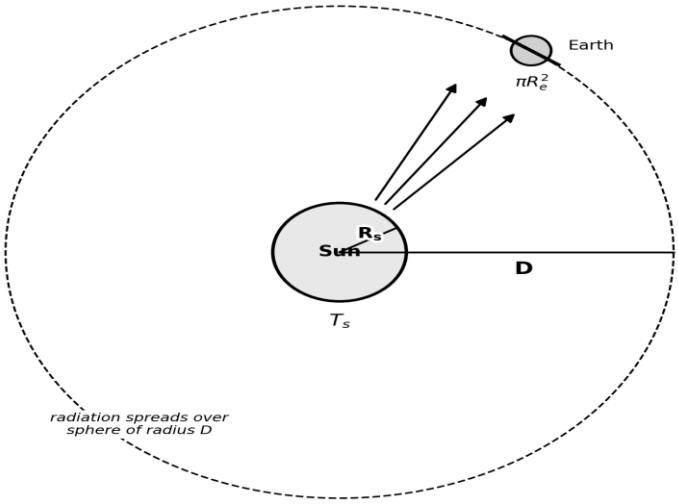


Figure: *Sun-Earth geometry for the solar-constant derivation. The Sun (radius R_s, surface temperature T_s) radiates uniformly in all directions. At Earth's orbital distance D, the radiation has spread over a sphere of area $4\pi D^2$. The Earth intercepts a cross-sectional area πR_e^2 of this radiation.*

Deriving the solar constant

Treat the Sun as a blackbody of radius R_s and surface temperature T_s. By Stefan-Boltzmann, the total power radiated by the Sun is:

$$P_{Sun} = \sigma(4\pi {R_s}^2)T_s^4$$

This radiation spreads uniformly in all directions. At a distance D from the Sun (the Earth's orbital radius), the same total power is distributed over a sphere of surface area $4\pi D^2$. The power per unit area at distance D is therefore:

$$\varphi_s = \frac{P_{Sun}}{4\pi D^2} = \frac{\sigma(4\pi {R_s}^2)T_s^4}{4\pi D^2}$$

The factors of 4π cancel, leaving:

$$\boldsymbol{\varphi_s = \sigma\left(\frac{R_s}{D}\right)^2 T_s^4}$$

This quantity φ_s is the **solar constant**: *the power per unit area of solar radiation arriving at the top of Earth's atmosphere, measured perpendicular to the Sun's rays.* Its unit is Wm^{-2}.

Earth's equilibrium temperature

The Earth intercepts solar radiation over a circular cross-section of area $\pi {R_e}^2$ (the shadow the Earth would cast). The total power received by Earth is therefore:

$$P_{received} = \varphi_s \pi R_e^2$$

The Earth re-radiates this energy back into space from its entire surface area $4\pi {R_e}^2$ (treating the Earth as a blackbody sphere). At equilibrium, the power received equals the power emitted:

$$\varphi_s \pi R_e^2 = \sigma(4\pi R_e^2)T_e^4$$

The $\pi {R_e}^2$ cancels from both sides:

$$\varphi_s = 4\sigma T_e^4$$

Substituting the expression for φ_s:

$$\sigma\left(\frac{R_s}{D}\right)^2 T_s^4 = 4\sigma T_e^4$$

Cancelling σ and solving for T_e:

$$T_e^4 = \frac{{R_s}^2 T_s^4}{4D^2}$$

$$\boldsymbol{T_e = T_s\sqrt{\left(\frac{R_s}{2D}\right)}}$$

This is the equilibrium temperature of a planet at orbital distance D from a star of radius R_s and surface temperature T_s, assuming the planet is a blackbody with no atmosphere.

Numerical computation

Given:
Sun's radius $R_s = 6.96 \times 10^8$ m
Earth's orbital distance $D = 1.50 \times 10^{11}$ m
Sun's surface temperature $T_s = 5800K$

$$T_e = 5800K \times \sqrt{\frac{6.96 \times 10^8 \text{ m}}{2 \times 1.50 \times 10^{11} \text{ m}}} = 279K$$

Earth's equilibrium temperature without an atmosphere would be about 279K, which is 6°C. The observed mean surface temperature of Earth is about 288K (15°C). The difference of 9K is the **greenhouse effect**: the atmosphere traps outgoing infrared radiation and warms the surface above what bare radiative balance would predict. Without the greenhouse effect, Earth would be a frozen planet barely above the melting point of ice.

Understanding the role of reflection: albedo

The derivation above assumed that Earth absorbs every photon of sunlight that strikes it. In reality, a significant fraction of the incoming solar radiation is reflected straight back into space before it can be absorbed. Clouds reflect sunlight from their white upper surfaces. Ice caps and snow-covered land reflect it from theirs. Even the ocean reflects a portion at low angles. The fraction of incoming radiation that a planet reflects is called its **albedo**, given the symbol α (the Greek letter alpha). An albedo of $\alpha = 0$ means the planet absorbs everything (a perfect blackbody); an albedo of $\alpha = 1$ means the planet reflects everything (a perfect mirror). Earth's average albedo is approximately $\alpha \approx 0.30$, meaning about 30% of the sunlight arriving at the top of the atmosphere is reflected and only 70% is absorbed.

When the albedo is taken into account, the power actually absorbed by Earth is not $\varphi_s \pi {R_e}^2$ but $(1 - \alpha)\varphi_s \pi {R_e}^2$. The radiative equilibrium equation becomes:

$$(1 - \alpha)\varphi_s \pi {R_e}^2 = \sigma\left(4\pi {R_e}^2\right){T_e}^4$$

Solving for T_e gives a corrected equilibrium temperature that is lower than the bare-blackbody value. With $\alpha = 0.30$ and $\varphi_s = 1380 Wm^{-2}$, the corrected equilibrium temperature comes out to approximately 255K (−18°C), which is 33°C below the observed mean surface temperature of 288K (15°C). This 33°C difference is the true greenhouse lift: it is the full warming that Earth's atmosphere provides by trapping outgoing infrared radiation. The idealized calculation without albedo (279K, giving a 9K greenhouse lift) understated the effect because it assumed the planet absorbed more sunlight than it actually does.

Newton's Law of Cooling as a Limiting Case of Stefan's Law

Earlier in this chapter, Newton's law of cooling was stated as an empirical observation: the rate of heat loss is proportional to the excess temperature $(\theta - \theta_s)$. We can now show that Newton's law follows directly from Stefan-Boltzmann's law when the excess temperature is small compared to the absolute temperature of the surroundings.

Starting with the net radiative power loss from Stefan-Boltzmann:

$$P_{net} = \varepsilon\sigma A(T^4 - T_s^4)$$

Factoring the difference of fourth powers.
Recall the algebraic identity: $a^4 - b^4 = (a - b)(a + b)(a^2 + b^2)$
Applying this with $a = T$ and $b = T_s$:

$$T^4 - T_s^4 = (T - T_s)(T + T_s)(T^2 + T_s^2)$$

Now applying the condition that the excess temperature is small: $(T - T_s) \ll T_s$.
This means $T \approx T_s$, so we can approximate the remaining factors:

$$(T + T_s) \approx 2T_s$$
$$(T^2 + T_s^2) \approx 2T_s^2$$

Substituting these approximations into the factored expression:

$$T^4 - T_s^4 \approx (T - T_s) \times 2T_s \times 2T_s^2 = 4T_s^3(T - T_s)$$

The net radiative power loss therefore becomes:

$$\mathbf{P_{net}} \approx \varepsilon\sigma A \times 4T_s^3(T - T_s) = 4\varepsilon\sigma A T_s^3(\mathbf{T - T_s})$$

This is exactly Newton's law of cooling. The rate of heat loss is directly proportional to the excess temperature, with a proportionality constant:

$$k = 4\varepsilon\sigma A T_s^3$$

So Newton's constant k is not truly a constant: it depends on the emissivity, the surface area, and the cube of the surroundings temperature. But for a given body cooling in a given room, all of these quantities are fixed, and k behaves as a constant throughout the cooling process.

This derivation reveals that Newton's law is a small-excess-temperature approximation to Stefan-Boltzmann's law, valid whenever $(T - T_s) \ll T_s$, which is the condition stated in the convection section.

With this solid theoretical foundation in place, we are now ready to see these ideas come alive through four more examples.

BINDER Example 19

Describe what happens to a blackbody as it is constantly heated, starting from room temperature. In your answer, describe the colour progression, the shift of the peak wavelength, and state which quantity increases most rapidly with temperature.

Solution

At room temperature (about 300K), the blackbody emits radiation entirely in the infrared, far beyond the range of human vision. The body appears dark to the eye, and its peak wavelength, given by Wien's law $\lambda_{max} = b/T$, is at about $2.90 \times 10^{-3}\,\text{mK} / 300\text{K} \approx 9700\text{nm}$, deep in the infrared.

As the temperature rises, the peak wavelength shifts to shorter values. At about 800K, the peak enters the far-red end of the visible spectrum and the body begins to glow dull red. Continued heating brings the colour progression through bright red, orange, and yellow. At about 5800K (the temperature of the Sun's surface), the peak is in the middle of the visible spectrum and the body appears approximately white, because significant radiation is emitted at all visible wavelengths simultaneously. Above about 10000K, the peak moves into the blue-violet range and the body appears blue-white.

Throughout this progression, two quantities are increasing:

(1) The total radiated power per unit area, which scales as T^4 (Stefan-Boltzmann).

(2) The peak spectral emissive power, which scales as T^5.

Of the two, the peak spectral emissive power increases most rapidly, because T^5 grows faster than T^4. Doubling the temperature increases the total radiated power by a factor of 16 but increases the peak height by a factor of 32.

Making Sense of the Answer: *The colour progression is a direct visual demonstration of Wien's law: as the temperature rises, the peak moves to shorter wavelength, and the eye perceives whichever visible wavelength happens to be closest to the peak. The fact that the body eventually appears white (not a single colour) reflects the breadth of the blackbody curve: at high temperatures the curve is broad enough to span the entire visible spectrum, and the eye perceives the mixture as white.*

Think Like a Physicist: *When an exam question asks "what happens to a blackbody as it is heated?", three things must appear in the answer: (1) the peak shifts to shorter wavelength (Wien), (2) the total radiated power grows as the fourth power of absolute temperature (Stefan-Boltzmann), and (3) the peak height grows as the fifth power. State all three explicitly, because each earns a separate mark. Stating only the colour change without connecting it to the underlying laws is a description, not physics.*

REAL Example 20

A premature baby in Mbulu district hospital lies in an incubator. The baby's skin temperature is 36°C, skin surface area is 0.12m², and skin emissivity is 0.80. The air inside the incubator is maintained at a comfortable 35°C. However, the incubator has a thin plastic wall facing a window, and the inner surface of this wall is only 18°C. (a) Calculate the net rate of radiative heat loss from the baby toward the cold wall. (b) A senior nurse drapes a warm cloth over the plastic wall, raising its inner surface temperature to 33°C. Calculate the new net rate of radiative heat loss. (c) Using Prevost's theory, explain why warm air alone was not enough to keep the baby safe.

Solution

This is a Prevost problem made quantitative. The baby emits radiation and absorbs radiation from the enclosure walls simultaneously. When the wall is cold, the absorption is small and the net loss is large, even though the air is warm.

Given:

Baby's skin: $T = 36°\text{C} = 309\text{K}$, $A = 0.12\text{m}^2$, $\varepsilon = 0.80$

Cold wall: $T_s = 18°\text{C} = 291\text{K}$

Warm wall (after blanket): $T_s = 33°\text{C} = 306\text{K}$

Stefan-Boltzmann constant $\sigma = 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4}$

(a) Net radiative heat loss toward the cold wall at 291K:

$$P_{net} = \varepsilon\sigma A\left(T^4 - {T_s}^4\right) = 0.80 \times 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4} \times 0.12\text{m}^2 \times ((309\text{K})^4 - (291\text{K})^4) = 10.6\text{W}$$

The baby loses about 10.6W by radiation toward the cold wall.

(b) Net radiative heat loss after the blanket raises the wall to 306K:

$$P_{net} = 0.80 \times 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4} \times 0.12\text{m}^2 \times ((309\text{K})^4 - (306\text{K})^4) = 1.96\text{W}$$

After covering the wall, the baby loses only 1.96W by radiation.

(c) Prevost's theory says the baby emits and absorbs radiation simultaneously and continuously. The air temperature controls convective heat transfer, but it does not control radiative heat transfer: radiation passes straight through the air between the baby's skin and the enclosure walls without being absorbed by it. What matters for radiation is the temperature of the surfaces

that the baby is exposed to. When one of those surfaces is a cold window wall at 18°C, the wall sends back far less radiation than the baby emits toward it. The result is a large net radiative loss (10.6W) that the warm air cannot compensate. The air may be at 35°C, but the baby's skin is exchanging radiation with an 18°C wall, not with the air. This is why warm air alone is not enough: radiation is a separate heat-transfer channel that bypasses the air entirely and responds only to the temperatures of the solid surfaces surrounding the baby.

Making Sense of the Answer: *Covering the wall reduced the net radiative loss from 10.6W to about 2.0W, a factor of more than 5. For a tiny premature baby with a resting metabolic rate of only a few watts, this difference can be the margin between stable body temperature and dangerous hypothermia. The physics behind this is pure Prevost: raise the wall temperature and the wall sends more radiation back to the baby, increasing the absorbed power while the emitted power stays the same. The net loss shrinks because the absorption has risen, not because the emission has fallen.*

Think Like a Physicist: *Whenever a radiation problem involves a body inside an enclosure, check the temperature of the walls, not the temperature of the air. Air is nearly transparent to infrared radiation, so it plays almost no role in radiative heat exchange. The surfaces the body "sees" are what determine the absorbed power. A cold wall, a cold window, or a cold ceiling can drain a body by radiation even when the air is warm. This is the principle behind radiant heating panels, double-walled incubators, and the reason hospitals keep premature babies away from windows.*

HOT Example 21

The solar constant at Earth's orbit is measured to be $\varphi_s = 1370\text{Wm}^{-2}$. Earth orbits the Sun at a distance $D = 1.50 \times 10^{11}$ m, and the Sun's radius is $R_s = 6.96 \times 10^8$ m. (a) Using the solar constant formula, calculate the Sun's surface temperature. (b) The Sun's peak emission wavelength is measured to be about 500nm. Using Wien's law, independently calculate the Sun's surface temperature. (c) Comment on the agreement between the two methods and what it tells us about the Sun as a radiator.

Solution

(a) From the solar constant formula, rearranging for the surface temperature:

$$\varphi_s = \sigma\left(\frac{R_s}{D}\right)^2 T_s^4$$

Solving for T_s:

$$T_s^4 = \frac{\varphi_s}{\sigma\left(\frac{R_s}{D}\right)^2}$$

$$T_s = \sqrt[4]{\frac{\varphi_s}{\sigma\left(\frac{R_s}{D}\right)^2}} = \sqrt[4]{\frac{1370\text{Wm}^{-2}.}{5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4}\left(\frac{6.96 \times 10^8\,\text{m}}{1.50 \times 10^{11}\,\text{m}}\right)^2}} = 5800\text{K}$$

(b) From Wien's displacement law:

$$T_s = \frac{b}{\lambda_{max}} = \frac{2.90 \times 10^{-3}\,\text{mK}}{500 \times 10^{-9}\,\text{m}} = 5800\text{K}$$

(c) The two methods give exactly the same value 5800K. The agreement confirms that the Sun's radiation closely follows the Planck blackbody spectrum: the same temperature that explains the total radiated power also explains the wavelength at which the radiation peaks. If the Sun were not a good blackbody, these two temperatures would disagree, and the size of the disagreement would measure how far the Sun's spectrum departs from the Planck curve.

Making Sense of the Answer: *This cross-check is how astrophysicists validate stellar temperature measurements. A single temperature that simultaneously satisfies both Stefan-Boltzmann (total power) and Wien (spectral shape) is strong evidence that the object is radiating as a blackbody. For the Sun, the match is excellent. For other stars with thick atmospheres or unusual compositions, the two temperatures sometimes disagree, signalling that the simple blackbody model needs refinement for that particular star.*

Think Like a Physicist: *Whenever two independent methods give the same answer, the result is more trustworthy than either method alone. In experimental physics this is called a consistency check: if two different physical laws, applied to two different measurements, converge on the same number, then either both are correct or both are wrong in exactly the same way, which is vastly less likely. Building consistency checks into your problem-solving routine is one of the most powerful habits a physicist can develop.*

HOT Example 22

A blackbody sphere of radius 5.0cm sits in a room at $T_s = 300$K. (a) Using the expression for Newton's cooling constant derived from Stefan's law, calculate the value of k for this sphere. (b) The sphere is heated to 320K and allowed to cool. Verify that the small-excess condition is satisfied. (c) Using Newton's law with your calculated k, estimate the initial net rate

of heat loss. (d) Calculate the exact net rate of heat loss using the full Stefan-Boltzmann formula and comment on the accuracy of Newton's approximation.

Solution

Given:
Blackbody sphere: $\varepsilon = 1$, radius $r = 5.0\text{cm} = 0.050\text{m}$, so $A = 4\pi r^2 = 4\pi(0.050\text{m})^2 = 0.0314\text{m}^2$
Surroundings: $T_s = 300\text{K}$. Heated temperature: $T = 320\text{K}$
(a) Newton's cooling constant from Stefan's law:

$$k = 4\varepsilon\sigma A T_s^{\,3} = 4 \times 1 \times 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4} \times 0.0314\text{m}^2 \times (300\text{K})^3 = 0.192\text{WK}^{-1}$$

(b) The excess temperature is $T - T_s = 320\text{K} - 300\text{K} = 20\text{K}$. For Newton's law (as a limiting case of Stefan's law) to be valid, we need $(T - T_s) \ll T_s$. The ratio is $20\text{K} / 300\text{K} = 0.067$, which is about 7%. This is comfortably small (well below the rule-of-thumb threshold of about 10%), so the small-excess condition is satisfied and Newton's law should be a good approximation.
(c) Newton's law estimate of the initial net rate of heat loss:

$$P_{\text{Newton}} = k(T - T_s) = 0.192\text{WK}^{-1} \times 20\text{K} = 3.84\text{W}$$

(d) Exact Stefan-Boltzmann net formula:

$$P_{\text{exact}} = \varepsilon\sigma A\left(T^4 - T_s^{\,4}\right) = 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4} \times 0.0314\text{m}^2 \times ((320\text{K})^4 - (300\text{K})^4) = 4.25\text{W}$$

Comparing the two results: Newton's law gives 3.84W, the exact formula gives 4.25W. The Newton estimate is about 10% below the exact value. This is a reasonable approximation for an excess of 20K out of 300K. The error arises because the approximation $T \approx T_s$ is not perfect: the true fourth-power curve is slightly steeper than the linear tangent that Newton's law represents. For smaller excesses the agreement would be even closer; for larger excesses it would deteriorate.

Making Sense of the Answer: *The 10% discrepancy is itself a useful number. It tells the student that Newton's law is a practical tool for small excesses (a 7% excess gives a 10% error in the heat-loss rate), but not a perfect one. The exact Stefan-Boltzmann formula is always more accurate, and the price of using Newton's law is this systematic undershoot.*

Think Like a Physicist: *This problem demonstrates the physicist's habit of checking an approximation against the exact answer. Deriving an approximation is step one; computing the error it introduces is step two. An approximation without an error estimate is a guess, not physics. Whenever you use Newton's law of cooling, you should be able to state the condition under which it is valid and estimate how large the error is for the specific excess temperature in the problem.*

With thermal radiation now complete, including its three major laws (Stefan-Boltzmann, Wien, and Newton's law as a limiting case) and the theory of heat exchange (Prevost), we have covered all three heat-transfer mechanisms: conduction, convection, and radiation. The next section returns to Newton's law of cooling and takes up its mathematical consequences: the exponential cooling curve, the integration of the cooling equation, the half-cooling time, and the experimental methods used to verify the law.

NEWTON'S LAW OF COOLING

In the convection section we stated Newton's law of cooling as an empirical observation: the rate of heat loss from a hot body is proportional to the excess temperature of the body over its surroundings. In the radiation section we showed that this law is a small-excess-temperature approximation to Stefan-Boltzmann's law. This section takes Newton's law and does the mathematics: we derive the differential equation, integrate it to find how the temperature changes with time, extract the half-cooling time, and develop the two-interval method that is the most heavily tested numerical technique in heat-related examination questions.

From Heat Loss to Temperature Change

Newton's law of cooling, in its heat-loss form, says:

$$\frac{dQ}{dt} = k(\theta - \theta_s)$$

where $\frac{dQ}{dt}$ is the rate at which the body loses heat (in watts), k is a positive constant that depends on the body's surface properties and surroundings, θ is the body's temperature, and θ_s is the surroundings temperature. But in practice we do not measure heat loss directly; we measure temperature. We need to convert this law into one that tells us how fast the temperature falls.
When a body of mass m and specific heat capacity c loses a small amount of heat dQ, its temperature drops by $d\theta$, where:

$$dQ = -mc\,d\theta$$

The minus sign is present because dQ is positive (heat leaves the body) while $d\theta$ is negative (the temperature falls). Dividing both sides by dt:

$$\frac{dQ}{dt} = -mc\,\frac{d\theta}{dt}$$

Setting this equal to Newton's law:

$$-\mathrm{mc}\,\frac{\mathrm{d}\theta}{\mathrm{dt}} = \mathrm{k}(\theta - \theta_s)$$

Dividing both sides by $-\mathrm{mc}$:

$$\frac{\mathrm{d}\theta}{\mathrm{dt}} = -\frac{\mathrm{k}}{\mathrm{mc}}(\theta - \theta_s)$$

The quantity $\frac{\mathrm{k}}{\mathrm{mc}}$ is a positive constant for a given body cooling in a given environment. We give it the symbol λ (lambda):

$$\boldsymbol{\lambda = \frac{k}{mc}}$$

The cooling equation becomes:

$$\boldsymbol{\frac{d\theta}{dt} = -\lambda(\theta - \theta_s)}$$

The result is **Newton's law of cooling in *temperature form*** which can be stated as: *under forced convection (at any excess temperature) or natural convection with a small excess temperature (below about 30°C), the rate of fall of temperature of a body is proportional to the excess temperature of the body over its surroundings.*

The constant $\boldsymbol{\lambda}$ has units of inverse time (s^{-1} or min^{-1}) and is called the **cooling constant**. A large λ means rapid cooling; a small λ means slow cooling.

Integrating the Cooling Equation

The cooling equation is a separable first-order differential equation. To solve it, separate the variables and integrate.

$$\frac{\mathrm{d}\theta}{\theta - \theta_s} = -\lambda\,\mathrm{dt}$$

Integrate both sides. At $\mathrm{t} = 0$ the body's temperature is θ_i (the initial temperature). At time t the temperature is θ:

$$\int_{\theta_i}^{\theta} \frac{\mathrm{d}\theta}{\theta - \theta_s} = -\lambda \int_0^{t} \mathrm{dt}$$

$$[\ln(\theta - \theta_s)]_{\theta_i}^{\theta} = -\lambda[t]_0^t$$

$$\ln(\theta - \theta_s) - \ln(\theta_i - \theta_s) = -\lambda t$$

$$\boldsymbol{\ln\left(\frac{\theta - \theta_s}{\theta_i - \theta_s}\right) = -\lambda t}$$

This is the logarithmic form of Newton's law of cooling.

Taking the exponential of both sides:

$$\frac{\theta - \theta_s}{\theta_i - \theta_s} = \mathrm{e}^{-\lambda t}$$

Rearranging to make the temperature the subject:

$$\boldsymbol{\theta = \theta_s + (\theta_i - \theta_s)e^{-\lambda t}}$$

This is the **integrated form of Newton's law of cooling**. It says that *the temperature of a cooling body decays exponentially from its initial value θ_i toward the surroundings temperature θ_s, with a time constant $1/\lambda$.* The excess temperature $(\theta - \theta_s)$ decays by the same fraction in each equal time interval, which is the hallmark of exponential decay.

The Cooling Curve

The graph of temperature against time is called the cooling curve. Its shape follows directly from the integrated form.

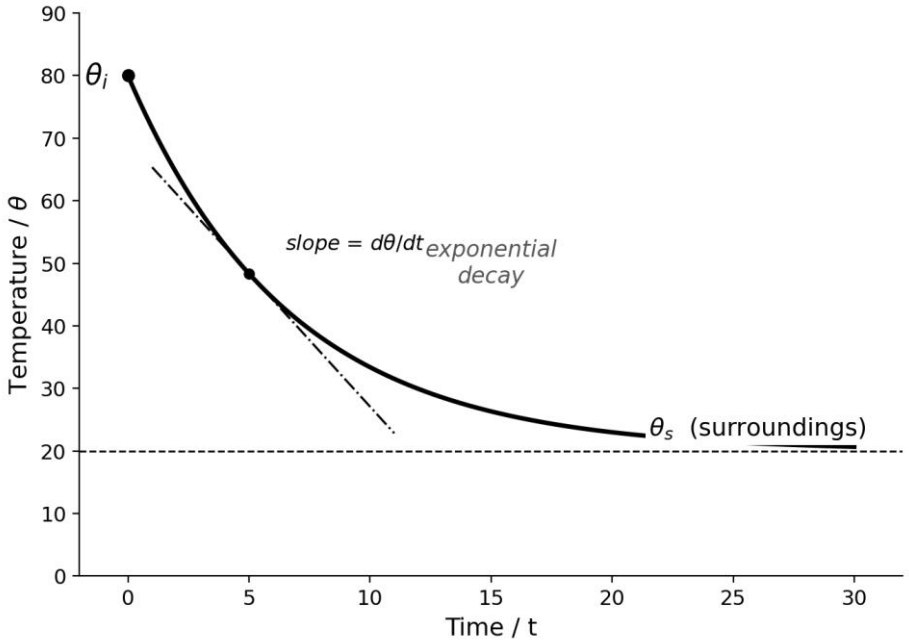


Figure: ***The cooling curve:*** *temperature θ plotted against time t. The curve starts at the initial temperature θ_i and decays exponentially toward the surroundings temperature θ_s, which appears as a horizontal asymptote. The tangent at any point gives the instantaneous rate of cooling dθ/dt at that moment.*

Key features of the cooling curve

1. The curve is steepest at the start (when the excess temperature is greatest) and becomes progressively flatter as the body approaches the surroundings temperature. This is because the rate of cooling is proportional to the excess: a large excess drives fast cooling, and fast cooling quickly reduces the excess, which slows the cooling further.
2. The curve has a horizontal asymptote at θ_s. Mathematically, the exponential factor $e^{-\lambda t}$ approaches zero as $t \to \infty$, so $\theta \to \theta_s$. But the exponential never reaches zero: the body gets closer and closer to θ_s without ever arriving. In practice, the body reaches a temperature indistinguishable from θ_s after a few multiples of $1/\lambda$.
3. The tangent to the curve at any point gives the instantaneous rate of cooling $\frac{d\theta}{dt}$ at that moment. The tangent is steep at the start and flat near the asymptote.

The Rate-of-Cooling Graph

If we plot the rate of cooling $-\frac{d\theta}{dt}$ against the excess temperature $(\theta - \theta_s)$, Newton's law predicts a straight line through the origin with gradient λ.

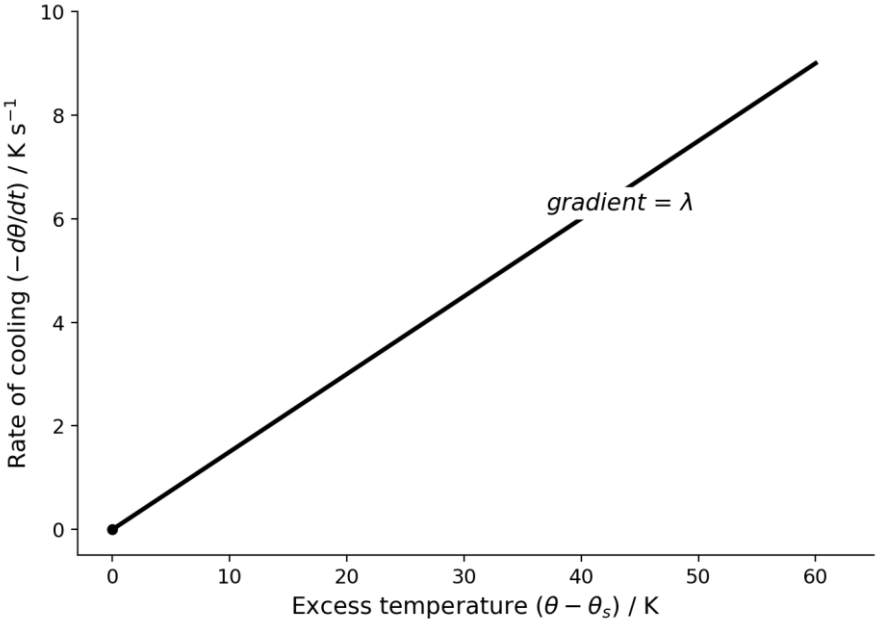


Figure: ***The rate-of-cooling graph.*** *A plot of the rate of temperature fall (−dθ/dt) against the excess temperature $(\theta - \theta_s)$ gives a straight line through the origin. The gradient of this line is the cooling constant λ.*

This graph is used experimentally to verify that a cooling body obeys Newton's law. If the data points fall on a straight line through the origin, the law is confirmed; if they curve, the cooling departs from the Newton regime (perhaps the excess is too large, and Dulong-Petit applies instead). The gradient of the line gives the numerical value of the cooling constant.

Half-Cooling Time

Just as radioactive decay has a half-life, Newton's cooling has a **half-cooling time**: *the time required for the excess temperature to fall to half its initial value.*

At time $t_{1/2}$, the excess temperature has fallen to half its initial value:

$$\theta - \theta_s = \frac{1}{2}(\theta_i - \theta_s)$$

Substituting into the logarithmic form:

$$\ln\left(\frac{\frac{1}{2}(\theta_i - \theta_s)}{\theta_i - \theta_s}\right) = -\lambda t_{1/2}$$

$$\ln\left(\frac{1}{2}\right) = -\lambda t_{1/2}$$

$$-\ln 2 = -\lambda t_{1/2}$$

$$\mathbf{t_{1/2} = \frac{\ln 2}{\lambda} = \frac{0.693}{\lambda}}$$

The half-cooling time depends only on λ, not on the initial temperature or the surroundings temperature. This is exactly analogous to the half-life of radioactive decay ($t_{1/2} = \ln 2/\lambda$), and for the same reason: exponential decay always takes the same time to halve, regardless of where you start measuring.

The Two-Interval Method

The most common type of Newton's-law problem in the examinations gives three temperatures (or two temperature drops) at known time intervals and asks the student to find a fourth (or third) temperature or a cooling time. The systematic way to solve such problems is the two-interval method.

The key idea is this: the logarithmic form of Newton's law applies to any two moments in time, not just from the start. If the body cools from temperature $\boldsymbol{\theta_1}$ (higher temperature) to temperature $\boldsymbol{\theta_2}$ (lower temperature) in time interval t, then:

$$\mathbf{\ln\left(\frac{\theta_1 - \theta_s}{\theta_2 - \theta_s}\right) = \lambda t}$$

If we know two such intervals (or one interval and a relationship between the two), we can write two equations and solve for the unknowns.

The equal-time shortcut

If both intervals have the **same** duration t, the right-hand sides of the two logarithmic equations are both λt, and we can set them equal or divide them. This produces a particularly clean ratio that eliminates λ entirely:

$$\mathbf{\frac{\theta_1 - \theta_s}{\theta_2 - \theta_s} = \frac{\theta_2 - \theta_s}{\theta_3 - \theta_s}}$$

where θ_1, θ_2, θ_3 are the temperatures at the start of the first interval, the end of the first interval (= start of the second), and the end of the second interval. This ratio says that the excess temperatures at three equally-spaced times form a geometric progression. It is the single most useful equation in Newton's-law problems.

In practical terms: if you know any three of the four quantities (the three temperatures and the surroundings temperature), you can find the fourth.

For now, let us pause the theory and bring the ideas down to earth with six carefully chosen worked examples.

A note on the regime of validity. *The problems that follow may use Newton's law at excess temperatures of 50K and above. Strictly speaking, the convection section showed that Newton's law is a small-excess approximation, valid when the excess is below about 30°C under natural convection. However, NECTA examination questions routinely apply Newton's law at much larger excesses without comment, and students are expected to use it unless the question specifically invokes Dulong-Petit. In practice, the two-interval method and the exponential decay formula work mathematically at any excess; the physical accuracy decreases at large excesses, but the examination expects the method regardless. When in doubt, use Newton's law unless the question tells you otherwise.*

BINDER Example 23

A body cools from 60°C to 40°C in 7 minutes. The surroundings are at 10°C. Find the temperature of the body after the next 7 minutes.

Solution

Given:

$\theta_1 = 60°C, \theta_2 = 40°C, \theta_s = 10°C, \theta_3 = ?$

The excess temperatures at equally-spaced times (of 7 minutes) form a geometric progression:

$$\frac{\theta_1 - \theta_s}{\theta_2 - \theta_s} = \frac{\theta_2 - \theta_s}{\theta_3 - \theta_s}$$

Substituting the excess temperatures:

$$\frac{60°C - 10°C}{40°C - 10°C} = \frac{40°C - 10°C}{\theta_3 - 10°C}$$

Solving for θ_3:

$$\theta_3 = 28°C$$

After the next 7 minutes, the body will be at 28°C.

Making Sense of the Answer: *The body cooled by 20°C in the first 7 minutes (from 60 to 40) but only by 12°C in the next 7 minutes (from 40 to 28). The rate of cooling has slowed because the excess temperature has shrunk: at the start the excess was 50°C, driving rapid cooling; after 7 minutes it was only 30°C, driving slower cooling; after 14 minutes it will be only 18°C. This is the signature of exponential decay: the temperature drop in each successive equal interval is smaller than the previous one.*

Think Like a Physicist: *In a two-interval problem with **equal time intervals**, always write the geometric-progression equation first. It eliminates λ and gives a direct relationship between the three temperatures and the surroundings temperature.*

BINDER Example 24

A body cools from 80°C to 50°C in 5 minutes. The surroundings are at 20°C. How long does it take to cool from 60°C to 30°C?

Solution

This problem gives two cooling intervals with different start and end temperatures but the same λ. Use the logarithmic form on each interval.

Interval 1: 80°C to 50°C in 5 minutes.

$$\ln\left(\frac{80°C - 20°C}{50°C - 20°C}\right) = \lambda \times 5\text{min}$$

$$\lambda = \frac{\ln 2}{5\text{min}} = 0.1386\ \text{min}^{-1}$$

Interval 2: 60°C to 30°C in time t.

$$\ln\left(\frac{60°C - 20°C}{30°C - 20°C}\right) = 0.1386\ \text{min}^{-1} \times t$$

$$t = 10.0\ \text{min}$$

It takes 10.0 minutes to cool from 60°C to 30°C.

Making Sense of the Answer: *The first interval (80 to 50, a 30°C drop) takes 5 minutes. The second interval (60 to 30, also a 30°C drop) takes 10 minutes, twice as long. A student who assumed equal drops take equal time would have answered 5 minutes and been wrong. The second interval takes longer because the body starts at a lower excess temperature (40°C versus 60°C) and ends at an even lower one (10°C versus 30°C), so the cooling is slower throughout. Equal temperature drops do not take equal time; exponential decay guarantees that each successive equal drop takes longer than the last.*

Think Like a Physicist: *The trap in this problem is the temptation to say "same drop, same time". This is wrong for exponential cooling and right only for linear cooling. The cure is to always use the logarithmic equation, never a shortcut based on temperature drops alone.*

REAL Example 25

A body is discovered in a hotel room in **Moshi**. The room temperature is a steady 20°C. A police doctor measures the body temperature at 10:00 PM and finds it to be 30°C. One hour later, at 11:00 PM, the body temperature has fallen to 27°C. Normal body temperature is 37°C. Assuming Newton's law of cooling applies, estimate the time of death.

Solution

The strategy: use the two measured temperatures (30°C and 27°C, one hour apart) to find the cooling constant λ. Then work backwards from the known starting temperature of 37°C to find how long ago the body was at that temperature.

Finding λ from the two readings:

At 10:00 PM: $\theta = 30°C$, excess = $\theta_1 - \theta_s = 30°C - 20°C = 10°C$

At 11:00 PM: $\theta = 27°C$, excess = $\theta_2 - \theta_s = 27°C - 20°C = 7°C$

$$\ln\left(\frac{10}{7}\right) = \lambda \times 1\text{hr}$$

$$\lambda = 0.3567\ \text{hr}^{-1}$$

Finding the time t from the moment of death to 10:00 PM:

At the moment of death, the excess was:

$$\theta_i - \theta_s = 37°C - 20°C = 17°C$$

At 10:00 PM, the excess was 10℃. The time elapsed from death to 10:00 PM is found from:

$$\ln\left(\frac{17}{10}\right) = \lambda t = 0.3567\ \text{hr}^{-1} \times t$$

$$t = 1.49\ \text{hr} \approx 1\ \text{hr}\ 29\ \text{min}$$

The body had been cooling for about 1 hour 29 minutes before the 10:00 PM reading.

Hence the estimated time of death is approximately **8:31 PM**.

Making Sense of the Answer: *The method relies on two assumptions: that Newton's law applies (the excess is small enough for the linear regime) and that the room temperature stayed constant. At 10:00 PM the excess was 10℃, comfortably within Newton's regime. At the moment of death the excess was 17℃, still reasonable. In real forensic work, pathologists use more sophisticated models that account for clothing, body mass, and ambient conditions, but the basic principle is exactly this: measure two temperatures at known times, compute the cooling constant, and extrapolate backwards.*

Think Like a Physicist: *Working Newton's law backwards in time is the same logarithmic equation used forwards; the only difference is that the answer gives a time before the measurement rather than after it. Any problem that gives you two temperature readings at known times and asks "when was the body at a different temperature?" is a forensic-cooling problem in disguise, whether the object is a human body, a cup of coffee, or a piece of molten metal.*

REAL Example 26

Mr. Akilikubwa presents the following student's working on the board. The question was: *"A body cools from 84°C to 60°C in 8 minutes. Surroundings are at 12°C. Find the temperature after the next 8 minutes."* The student wrote: *"The body dropped 24°C in 8 minutes. By Newton's law, it will drop another 24°C in the next 8 minutes. Temperature = 60 − 24 = 36°C."* Mr. Akilikubwa turns to the class: *"This student has made a fundamental error."*

Where exactly is the mistake? Correct the working and find the right answer.

Solution

The student's mistake: assuming that equal temperature drops occur in equal time intervals. This is true only for linear cooling, which does not describe real cooling. Newton's law says the rate of cooling is proportional to the excess temperature, not constant. As the body cools, the excess shrinks, the driving force weakens, and each successive equal time interval produces a smaller drop than the previous one.

Correct working using the geometric-progression method:

Excess temperatures:

$$\theta_1 - \theta_s = 84°C - 12°C = 72°C$$
$$\theta_2 - \theta_s = 60°C - 12°C = 48°C$$

Geometric progression:

$$\frac{72°C}{48°C} = \frac{48°C}{\theta_3 - 12°C}$$

$$\theta_3 = 44°C$$

The correct answer is 44°C, not 36°C.

The student's error gave a temperature that is 8°C too low.

Making Sense of the Answer: *The student predicted a 24°C drop; the correct drop is only 16°C (from 60 to 44). In the first 8 minutes the body dropped 24°C (from 84 to 60) because the excess was large (72°C at the start). In the next 8 minutes the excess is smaller (48°C at the start), so the drop must be smaller. The ratio of consecutive drops (16/24 = 0.667) equals the ratio of consecutive excess temperatures (48/72 = 0.667). This constant ratio is the geometric-progression property: it is the fingerprint of exponential decay, and the surest sign that a student has understood Newton's law rather than guessed at it.*

Think Like a Physicist: *This error-correction format is a powerful study technique. Understanding why a wrong method gives the wrong answer teaches more than simply applying the right method mechanically. The student who can explain why equal drops do not take equal time has understood Newton's law at a deeper level than the student who merely plugs numbers into the formula. In an exam, if you are unsure of your method, check whether your answer implies equal drops in equal times. If it does, you have almost certainly made this error.*

HOT Example 27

A body cools from 70°C to 50°C in 10 minutes, and from 50°C to 35°C in the next 10 minutes. Find the temperature of the surroundings.

Solution

Two equal time intervals of 10 minutes. The geometric-progression equation applies, but this time the unknown is θ_s rather than a temperature or a time.

The excess temperatures at three equally-spaced times form a geometric progression:

$$\frac{\theta_1 - \theta_s}{\theta_2 - \theta_s} = \frac{\theta_2 - \theta_s}{\theta_3 - \theta_s}$$

Cross-multiplying:

$$(\theta_1 - \theta_s)(\theta_3 - \theta_s) = (\theta_2 - \theta_s)^2$$

Substituting $\theta_1 = 70°C, \theta_2 = 50°C, \theta_3 = 35°C$:

$$(70 - \theta_s)(35 - \theta_s) = (50 - \theta_s)^2$$

$$2450 - 105\theta_s + {\theta_s}^2 = 2500 - 100\theta_s + {\theta_s}^2$$

$$-50 = 5\theta_s$$

$$\theta_s = -10°C$$

The surroundings temperature is −10°C.

Making Sense of the Answer: *A negative surroundings temperature may look surprising for a body that starts at 70°C, but it is physically sensible: the body is cooling in a cold environment (perhaps outdoors on a chilly night or in a freezer room). The key test is that all three excess temperatures must be positive, and they are: 70 − (−10) = 80°C, 50 − (−10) = 60°C, 35 − (−10) = 45°C. Check the geometric progression: 80/60 = 1.333, and 60/45 = 1.333. They match. The answer is correct.*

Think Like a Physicist: *When the surroundings temperature is the unknown, the geometric-progression equation produces* ***a linear equation in θ_s*** *(because the quadratic terms cancel). This is a reliable algebraic pattern: learn it, and you can solve any "find the surroundings" problem in under a minute. The cancellation of the squared terms is not a coincidence; it is a consequence of the geometric-progression structure, and it happens* ***every time****.*

HOT Example 28

A hot object has a half-cooling time of 6 minutes in a room at 25°C. The object starts at 105°C. (a) Find the temperature after 6 minutes, after 12 minutes, and after 18 minutes. (b) How long does it take for the object to reach 30°C?

Solution

The half-cooling time halves the **excess temperature**, not the temperature itself. The initial excess is:

$$\theta_i - \theta_s = 105°C - 25°C = 80°C$$

(a) After each half-cooling time of 6 minutes, the excess halves:

After 6 min: excess = $\frac{80°C}{2} = 40°C$, so $\theta = 25°C + 40°C = 65°C$

After 12 min: excess = $\frac{40°C}{2} = 20°C$, so $\theta = 25°C + 20°C = 45°C$

After 18 min: excess = $\frac{20°C}{2} = 10°C$, so $\theta = 25°C + 10°C = 35°C$

(b) At 30°C, the excess is $30°C - 25°C = 5°C$. The excess must fall from 80°C to 5°C, which is a factor of $\frac{80}{5} = 16$. Since $2^4 = 16$, this requires 4 half-cooling times:

$$t = 4 \times 6\,\text{min} = \mathbf{24\ min}$$

Alternatively, using $\lambda = \frac{\ln 2}{6\,\text{min}} = 0.1155\ \text{min}^{-1}$:

$$t = \frac{\ln\left(\frac{80}{5}\right)}{0.1155\ \text{min}^{-1}} = \frac{\ln 16}{0.1155}\,\text{min} = \frac{2.773}{0.1155}\,\text{min} = 24.0\,\text{min}$$

The object reaches 30°C after 24 minutes.

Making Sense of the Answer: *A common mistake is to halve the temperature instead of the excess. A student who halves 105 gets 52.5, then 26.25, and concludes the body reaches 25°C after only two half-cooling times (12 minutes). The correct sequence (65, 45, 35, 30) shows the body is still at 35°C after 18 minutes and does not reach 30°C until 24 minutes. The error comes from forgetting that* ***Newton's law describes the decay of the excess over the surroundings, not the decay of the temperature itself.***

Think Like a Physicist: *Whenever a problem gives you the half-cooling time, convert it to the cooling constant immediately:* $\lambda = \frac{ln\,2}{t_{1/2}}$*. From there, every calculation uses the standard logarithmic equation. The half-cooling time is a convenient input format, not a substitute for the logarithmic method. For problems where the answer is an exact number of half-cooling times (as in part (a)), halving the excess is faster. For problems where the answer is not a whole number of half-cooling times (which is the more common case in exams),* ***the logarithmic equation is the only reliable method.***

Newton's law of cooling is now fully in our hands: stated, derived, integrated, graphed, and practised on the types of problem that examiners favour. The next section takes up the experimental determination of thermal conductivity, using Searle's apparatus for good conductors and Lee's disc for poor conductors, both of which rely on the steady-state and cooling-curve techniques we have built across this chapter.

EXPERIMENTAL DETERMINATION OF THERMAL CONDUCTIVITY

Everything we have derived so far in this chapter depends on knowing the thermal conductivity of the material in question. But how is that number measured in the first place? This section describes the two standard laboratory methods: **Searle's apparatus** *for good conductors* and **Lee's disc** *for poor conductors*. Both methods rest on the same principle but use different experimental strategies to suit the very different thermal behaviours of metals and insulators.

The Underlying Principle

In the steady state, Fourier's law gives the rate of heat flow through a uniform sample:

$$\frac{dQ}{dt} = \frac{kA(\theta_1 - \theta_2)}{x}$$

where k is the thermal conductivity, A is the cross-sectional area, θ_1 and θ_2 are the temperatures of the hot and cold faces, and x is the thickness of the sample. To measure k, we must independently measure all other quantities in this equation: the rate of heat flow $\frac{dQ}{dt}$, the area A, the thickness x, and the two surface temperatures. The two methods differ only in how they measure $\frac{dQ}{dt}$.

Searle's Apparatus (Good Conductors)

Searle's method is designed for materials with high thermal conductivity, such as copper, brass, aluminium, and iron. The sample is a long metal bar, typically 30cm to 50cm, with a large cross-sectional area.

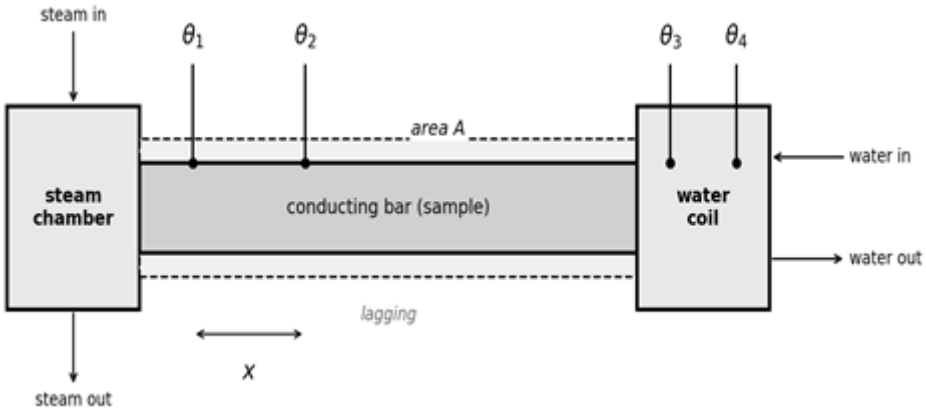


Figure: ***Searle's apparatus****. Steam enters the chamber at the left, maintaining the hot end of the bar at a steady temperature. The bar is lagged to prevent heat loss from its sides. Two thermometers (θ₁, θ₂) measure the temperature at two points along the bar, separated by a known distance x. At the cold end, water flows through a coil, entering at θ₃ and leaving at θ₄.*

How it works

Steam is passed continuously through the chamber at the hot end of the bar, maintaining a steady temperature. At the cold end, cold water flows through a copper coil at a known, measured rate. The bar is completely lagged along its length to prevent heat escaping from its sides. Under these conditions, all the heat that enters the bar at the hot end must exit through the cold end into the water. When the four thermometer readings become constant, the apparatus has reached its steady state.

At steady state, the rate of heat conducted through the bar equals the rate of heat absorbed by the water:

$$\frac{kA(\theta_1 - \theta_2)}{x} = \frac{m}{t}c(\theta_4 - \theta_3)$$

where $\frac{m}{t}$ is the mass of water flowing through the coil per unit time (measured by collecting and weighing water over a timed interval) and c is the specific heat capacity of water. The only unknown is k, so rearranging:

$$\mathbf{k = \frac{\frac{m}{t}\,c(\theta_4 - \theta_3)\,x}{A(\theta_1 - \theta_2)}}$$

Why the bar must be lagged

Without lagging, heat would escape from the sides of the bar into the surrounding air, and not all of the heat entering at the hot end would reach the cold end. The equation above assumes that the rate of heat flow is the same at every cross-section along the bar. If heat leaks from the sides, this assumption fails, and the measured value of k will be too low.

Why a long bar is needed for good conductors

A good conductor carries heat so efficiently that even a small temperature difference drives a large heat flow. If the bar were short, the temperature difference $(\theta_1 - \theta_2)$ between the two measurement points would be too small to read accurately on a

mercury thermometer. A long bar gives a larger separation between the thermometers and a larger, more measurable temperature difference.

Lee's Disc (Poor Conductors)

Lee's disc method is designed for materials with low thermal conductivity, such as cardboard, glass, rubber, cork, and wood. The sample is a thin disc, typically a few millimetres thick.

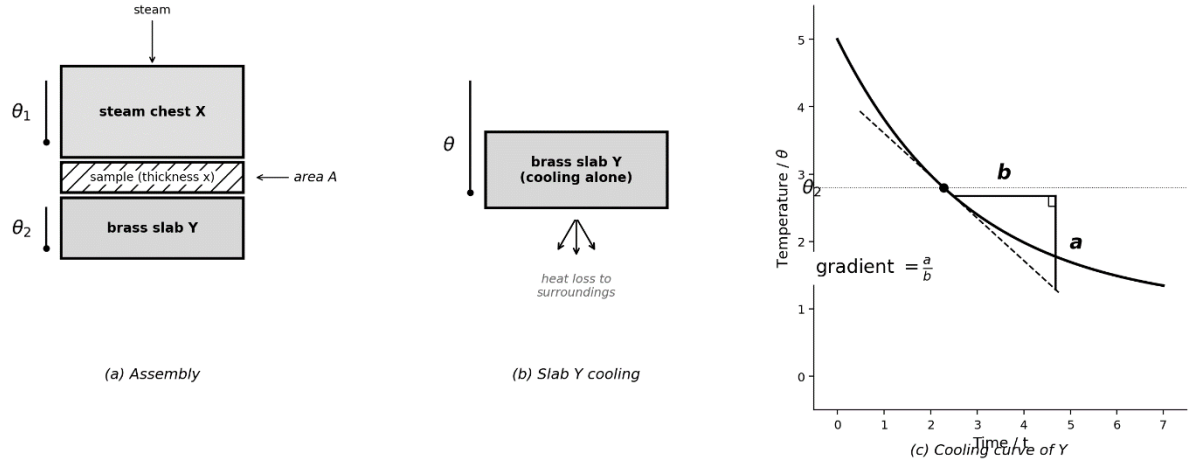


Figure: ***Lee's disc method.*** *(a) The assembly: the thin sample disc is sandwiched between the steam chest X (at θ_1) and the brass slab Y (at θ_2). (b) The brass slab Y is removed and allowed to cool alone, with a thermometer recording its temperature. (c) The cooling curve of Y: the tangent at the temperature θ_2 gives the gradient a/b, which is the rate of temperature fall at that temperature.*

How it works

The experiment has two stages.

Stage 1 (steady state): Steam is passed through the chest X, heating the top face of the sample disc. Heat conducts through the disc and into the brass slab Y below. The assembly is left until the thermometer readings θ_1 (on X) and θ_2 (on Y) become constant. At this point the steady state has been reached: the rate of heat conducted through the sample equals the rate at which slab Y loses heat to its surroundings.

Stage 2 (cooling curve): The sample disc and the steam chest are removed. The brass slab Y, now at temperature θ_2, is left to cool alone in the air. Its temperature is recorded at regular intervals and plotted as a cooling curve (θ vs t). A tangent is drawn to the curve at the point where $\theta = \theta_2$. The gradient of this tangent, measured as $\frac{a}{b}$ from the tangent triangle, gives the rate of temperature fall of slab Y at the steady-state temperature.

Deriving the working equation

At steady state, the rate of heat conducted through the sample equals the rate at which slab Y loses heat. The rate of heat loss by slab Y at temperature θ_2 is:

$$\frac{dQ}{dt} = mc\frac{a}{b}$$

where m is the mass of slab Y, c is its specific heat capacity, and $\frac{a}{b}$ is the magnitude of the gradient of the cooling curve at $\theta = \theta_2$. Setting this equal to the Fourier conduction rate through the disc:

$$\frac{kA(\theta_1 - \theta_2)}{x} = mc\frac{a}{b}$$

Solving for k:

$$\mathbf{k = \frac{mc\frac{a}{b}\,x}{A(\theta_1 - \theta_2)}}$$

Why a thin disc is needed for poor conductors

A poor conductor resists heat flow strongly. If the sample were thick, the cold face (slab Y) would take a very long time to warm up, and the apparatus might never reach a true steady state within a practical laboratory session. A thin disc ensures that the heat flow rate is large enough to bring slab Y to a measurable temperature within a reasonable time.

Why Different Methods for Different Conductors

The choice of method is not arbitrary; it is dictated by the physics of the material under test.

Good conductors need a **long** sample. A short bar of copper would have such a tiny temperature difference between its ends that the measurement would be swamped by thermometer uncertainty. A long bar gives a measurable $(\theta_1 - \theta_2)$. The water-coil calorimeter at the cold end measures the heat flow rate directly and accurately.

Poor conductors need a **thin** sample. A thick slab of cardboard would never reach steady state in a reasonable time. A thin disc reaches steady state quickly. The cooling-curve method replaces the water-coil calorimeter because the heat flow rate through a poor conductor is too small to heat water measurably; instead, it is measured indirectly through the rate of temperature fall of the brass slab.

Attempting to use Searle's method on cardboard would fail because the heat flow would be too small to warm the water. Attempting to use Lee's disc on a 1-metre iron bar would fail because the bar would never reach steady state in a thin-disc arrangement. Each method is tailored to its material.

With the ideas now simmering nicely, let us serve them properly through four worked examples and enjoy the flavour of physics in action.

BINDER Example 29

State the principle on which Searle's apparatus is based. Identify the role of each of the four thermometers and explain why the conducting bar must be lagged.

Solution

Principle:

In the steady state, the rate of heat conducted through the bar equals the rate of heat absorbed by water flowing through the coil at the cold end. Fourier's law gives the conduction rate; the water-flow calorimetry equation gives the absorption rate. Equating the two allows the thermal conductivity k to be calculated.

Role of thermometers:

- Thermometers θ_1 and θ_2 measure the temperature at two points along the bar, separated by a known distance x. Their readings give the temperature gradient along the bar, which appears in Fourier's law as $\frac{(\theta_1 - \theta_2)}{x}$.
- Thermometers θ_3 and θ_4 measure the inlet and outlet temperatures of the cooling water. Their difference $(\theta_4 - \theta_3)$ gives the temperature rise of the water, which, combined with the water flow rate and specific heat capacity, gives the rate of heat absorption.

Explanation on lagging bar:

The bar must be lagged to prevent heat loss from its sides. Without lagging, some of the heat entering at the hot end would escape to the surrounding air before reaching the cold end, and the water at the cold end would absorb less heat than the bar conducted. The calculated value of k would then be too low, because the equation assumes all conducted heat reaches the water.

Making Sense of the Answer: *Lagging is not a minor detail; it is the single most important experimental precaution in Searle's method. Without it, the fundamental equality (heat in = heat out) on which the entire method rests does not hold, and the experiment gives a systematically wrong answer.*

Think Like a Physicist: *In any experimental-method question, identify the equation being used and then explain each measurement in terms of a quantity in that equation. This maps every piece of apparatus to a physics symbol, which is exactly what the examiner wants to see. A description of the apparatus without linking it to the equation is a technician's account, not a physicist's.*

HOT Example 30

A copper rod 2.0m long and 1.0cm in radius is maintained at 250°C at one end. In the steady state, heat flows across each cross-section at a rate of 2.1Js⁻¹. The thermal conductivity of copper is $k = 380\text{Wm}^{-1}\text{K}^{-1}$. Find the temperature of the other end.

Solution

Fourier's law, rearranged to find the cold-end temperature.

Given:

Length $x = 2.0\text{m}$, radius $r = 1.0\text{cm} = 0.010\text{m}$

Area $A = \pi r^2 = \pi(0.010\text{m})^2 = 3.14 \times 10^{-4}\ \text{m}^2$

Hot-end temperature $\theta_1 = 250°\text{C}$, heat-flow rate $\frac{dQ}{dt} = 2.1\text{Js}^{-1} = 2.1\text{W}$

From Fourier's law:

$$\frac{dQ}{dt} = \frac{kA(\theta_1 - \theta_2)}{x}$$

Rearranging:

$$\theta_1 - \theta_2 = \frac{\frac{dQ}{dt} \times x}{kA} = \frac{2.1W \times 2.0m}{380Wm^{-1}K^{-1} \times 3.14 \times 10^{-4}\,m^2} = 35.2K = 35.2°C$$
$$\theta_2 = 250°C - 35.2°C = 214.8°C \approx 215°C$$

The cold end of the rod is at about 215°C.

Making Sense of the Answer: *A 2-metre copper rod carrying only 2.1 watts has a temperature drop of just 35°C across its entire length. This tiny gradient is the signature of a good conductor: copper transmits heat so readily that even a modest heat flow requires almost no temperature difference to drive it. This is also why Searle's method uses a long bar: a short bar of copper would give an even smaller temperature difference, making the measurement impossibly imprecise.*

Think Like a Physicist: *When a problem asks for the temperature at the other end of a conductor, rearrange Fourier's law for the temperature difference first, then subtract from the known temperature. Never rearrange for the unknown temperature directly; the algebra is cleaner when you isolate the difference.*

HOT Example 31

A brass boiler has a base of area 0.15m² and thickness 1.0cm. The boiler boils water at a rate of 6.0kg per minute. The thermal conductivity of brass is $k = 109Wm^{-1}K^{-1}$ and the specific latent heat of vapourisation of water is $L_v = 2256kJkg^{-1}$. Find the temperature of the surface in contact with the flame.

Solution

The boiler base conducts heat from the flame to the water. The water is boiling, so its temperature is 100°C. The rate of heat flow equals the rate at which the water absorbs latent heat to vapourise.

Given:
Area $A = 0.15m^2$, thickness $x = 1.0cm = 0.010m$
Water boiled: $\frac{m}{t} = 6.0kgmin^{-1} = 0.10kgs^{-1}$
Latent heat $L_v = 2256kJkg^{-1} = 2.256 \times 10^6\ Jkg^{-1}$
Rate of heat flow through the base:

$$\frac{dQ}{dt} = \frac{m}{t} \times L_v = 0.10kgs^{-1} \times 2.256 \times 10^6\ Jkg^{-1} = 2.256 \times 10^5\ W$$

From Fourier's law, the temperature difference across the base:

$$\theta_1 - \theta_2 = \frac{\frac{dQ}{dt} \times x}{kA} = \frac{2.256 \times 10^5\ W \times 0.010m}{109Wm^{-1}K^{-1} \times 0.15m^2} = 138K = 138°C$$
$$\theta_1 = 100°C + 138°C = 238°C$$

The flame-contact surface of the boiler base is at about 238°C.

Making Sense of the Answer: *The flame side of the boiler is 138°C hotter than the water side. This large temperature difference across just 1cm of brass reflects the enormous rate of heat flow: 225 kilowatts passing through 0.15m² of brass is about 1500kWm⁻², which is comparable to an industrial furnace. The reason the base does not melt is that brass has a melting point above 900°C, well above the 238°C flame-contact temperature. If the water were to boil dry, however, the heat flow would stop being absorbed by vapourisation, the base temperature would rise rapidly, and the base could warp or melt.*

Think Like a Physicist: *Boiling-rate problems are Fourier problems with a twist: the rate of heat flow is not given directly but must be computed from the mass of water boiled per second and the latent heat of vapourisation. Convert the boiling rate to watts first ($dQ/dt = dm/dt \times L_v$), then use Fourier's law in the normal way. The water-side temperature is always 100°C (at standard pressure).*

HOT Example 32

(a) Derive the working equation for the Lee's disc method of measuring the thermal conductivity of a poor conductor. (b) Explain why Searle's method cannot be used to measure the thermal conductivity of cardboard, and why Lee's disc method cannot be used for a 1-metre iron bar.

Solution

(a) In the steady state, the rate of heat conducted through the thin sample disc equals the rate at which the brass slab Y loses heat to its surroundings. By Fourier's law, the conduction rate through the disc of thickness x, area A, with face temperatures θ_1 (hot, from the steam chest) and θ_2 (cold, on slab Y) is:

$$\left(\frac{dQ}{dt}\right)_{conduction} = \frac{kA(\theta_1 - \theta_2)}{x}$$

The rate of heat loss by slab Y at temperature θ_2 is determined from the cooling curve. When slab Y cools alone, its temperature falls at a rate $\frac{d\theta}{dt}$ that can be read from the tangent to the cooling curve at $\theta = \theta_2$. The magnitude of this gradient is $\frac{a}{b}$ (from the tangent triangle). The rate of heat loss is therefore:

$$\frac{dQ}{dt_{loss}} = mc\frac{a}{b}$$

where m is the mass and c is the specific heat capacity of the brass slab. Equating the two rates:

$$\frac{kA(\theta_1 - \theta_2)}{x} = mc\frac{a}{b}$$

Solving for k:

$$k = \frac{mc\frac{a}{b}x}{A(\theta_1 - \theta_2)}$$

(b) Two reasons, one for each direction:

Searle's method cannot be used for cardboard because cardboard has very low thermal conductivity. The rate of heat flow through the sample would be so small that the cooling water at the cold end would barely warm up: the temperature rise $(\theta_4 - \theta_3)$ would be unmeasurably small. Without a measurable heat-flow rate, k cannot be calculated. Additionally, a long bar of cardboard would take an impractically long time to reach steady state.

Lee's disc method cannot be used for a 1-metre iron bar because the method requires a thin disc sample. A 1-metre bar is far too thick: the cold face would take an extremely long time to reach a steady temperature, and the cooling-curve stage assumes that the slab Y is the only significant thermal mass below the sample, which is not true when the sample itself is a massive iron bar. Furthermore, iron is a good conductor: it would easily drive a large heat flow into slab Y, overwhelming the cooling-curve measurement. Searle's method with a water-coil calorimeter is the correct approach for iron.

Making Sense of the Answer: *The two methods are complementary, not interchangeable. Each is optimised for a specific range of thermal conductivities. The dividing line is roughly around k = 1Wm⁻¹K⁻¹: materials above this value (metals, some dense ceramics) suit Searle; materials below it (wood, cardboard, rubber, cork, glass wool) suit Lee's disc. Knowing which method to use for which material is itself a piece of physics knowledge that the examiner tests.*

Think Like a Physicist: *When asked "why can't method X be used for material Y?", always answer in terms of a specific quantity in the working equation that becomes unmeasurably small or impractically large. For Searle on cardboard: the water temperature rise becomes unmeasurable. For Lee on iron: the sample is too thick to reach steady state. Vague answers like "it doesn't work" score no marks; specific answers tied to the equation score full marks.*

With the experimental methods for measuring thermal conductivity now complete, we have covered the full arc of Chapter 3: the three heat-transfer mechanisms (conduction, convection, radiation), their quantitative laws, Newton's law of cooling, and the laboratory techniques that put numbers on the thermal properties of materials. The final section brings everything together by examining real-world applications that combine all three mechanisms.

APPLICATIONS OF HEAT TRANSFER

No real-world object loses or gains heat by just one mechanism. A cup of tea on a table cools simultaneously by conduction through the table surface, by convection as warm air rises from its surface, and by radiation from its sides toward the cooler walls of the room. The engineering of everyday thermal devices is the art of identifying which mechanism dominates and either amplifying it (if you want heat to flow) or suppressing it (if you want heat to stay). This section applies the physics of the entire chapter to a set of familiar devices and natural phenomena.

1. The vacuum flask

The vacuum flask (thermos flask) is the most elegant thermal device in everyday life. It suppresses all three heat-transfer mechanisms simultaneously, using a different physical strategy for each.

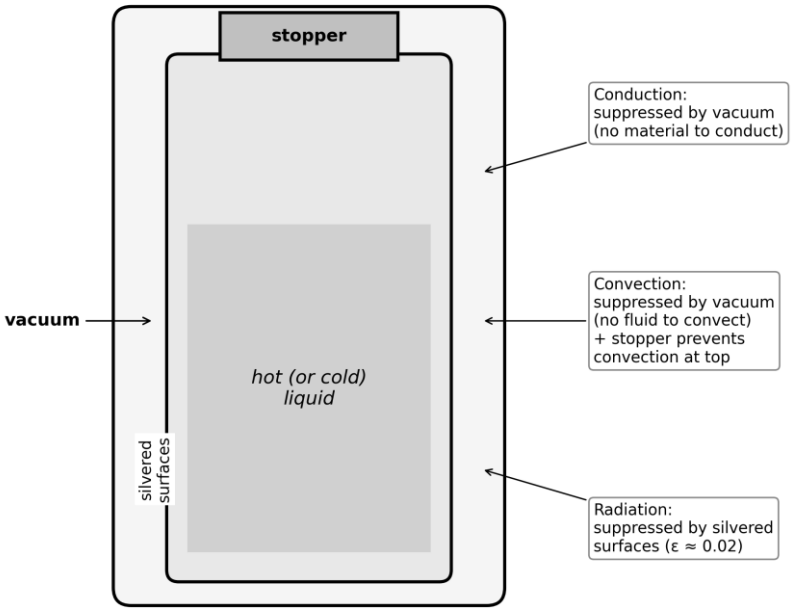


Figure: ***Cut-away view of a vacuum flask.*** *The vacuum between the double walls eliminates conduction and convection. The silvered inner surfaces reduce radiation to a minimum. The stopper prevents convection at the top.*

Conduction *is suppressed by the vacuum between the double walls.* With no material in the gap, there are no molecules to jostle their neighbours, and heat cannot conduct across. The only conduction paths remaining are through the glass walls themselves (which are thin) and through the stopper at the top.

Convection *is suppressed by the same vacuum.* With no air or other fluid in the gap, there is no medium to carry heat by bulk flow. The stopper at the top prevents warm air from rising out of the flask and being replaced by cooler air from the room.

Radiation *is suppressed by silvering the inner surfaces of the double walls.* Polished silver has an emissivity of only about 0.02, so the hot inner surface emits very little radiation outward, and the cold outer surface emits very little radiation inward. By Prevost's theory, low emissivity reduces both emission and absorption by the same factor, suppressing the net radiative exchange in both directions.

The result is a container that can keep a hot drink warm for many hours, or keep a cold drink cold for equally long. The flask does not "know" whether it is keeping heat in or keeping heat out; it simply reduces the rate of heat transfer in both directions by the same set of mechanisms.

2. The blanket

A blanket does not generate heat. It works by reducing the rate at which heat escapes from the body beneath it. The primary mechanism is **trapping air**: the loose weave of a blanket creates pockets of still air between the fibres, and air is one of the poorest thermal conductors. These trapped-air pockets *suppress both conduction* (air conducts poorly) and *convection* (the air pockets are too small for convection currents to develop). The fabric itself also *interrupts the direct radiation* path between the warm body and the cold room.

The same physics explains why a blanket can keep ice from melting. Wrapping ice in a blanket does not cool the ice; it reduces the rate at which heat from the warm room reaches the ice. The blanket suppresses the same three mechanisms in the same way, regardless of the direction of heat flow. This is a direct consequence of Fourier's law and Prevost's theory: anything that reduces the heat-transfer rate in one direction reduces it equally in the other.

3. Birds and animals in cold weather

When the temperature drops, birds fluff up their feathers and many animals curl into a ball. Both behaviours reduce heat loss, but by different physical strategies.

Fluffing feathers traps a thicker layer of still air close to the bird's body. Air is a poor conductor, so a thicker air layer increases the effective insulation thickness, raising the thermal resistance between the bird's warm skin and the cold surroundings. The trapped air also cannot convect because the feathers break it into small, isolated pockets. The bird has, in effect, put on a thicker blanket.

Curling into a ball reduces the surface area exposed to the cold environment. A sphere has the smallest surface-area-to-volume ratio of any shape, so by curling, an animal minimises the total area through which heat can escape by convection and radiation. It also reduces the area exposed to wind, suppressing forced convection. The physics is Stefan-Boltzmann (radiated power is proportional to surface area) and Newton's law of cooling (convective heat loss is proportional to surface area through the constant k, which depends on A).

4. The domestic "radiator"

The device commonly called a "radiator" in homes and offices is misnamed. It heats a room primarily by natural convection, not by radiation. Hot water flows through the metal fins of the radiator, heating the air in contact with them. The warm air rises (natural convection), travels across the ceiling, cools, descends along the far wall, and returns along the floor to the radiator, exactly as described in the convection section of this chapter. Radiation from the radiator surfaces does contribute, but it is a minor fraction of the total heat output because the radiator surface temperature (typically 50 to 70°C) is low enough that radiation is small compared to the vigorous convection current.

The name "radiator" is a historical holdover from the early days of central heating, when the distinction between convection and radiation was not yet well understood. In modern engineering, the same device is sometimes called a "convector", which is more accurate.

5. The greenhouse effect

The glass of a greenhouse is transparent to short-wavelength visible radiation from the Sun but opaque to the longer-wavelength infrared radiation emitted by the warm soil and plants inside. Sunlight passes through the glass and warms the interior. The warmed interior radiates infrared back toward the glass, but the glass absorbs it and re-radiates some of it back inward instead of letting it escape. The interior reaches a higher equilibrium temperature than it would without the glass, because the outgoing radiation is partially blocked.

Earth's atmosphere behaves in a similar way. Greenhouse gases (mainly water vapour and carbon dioxide) are transparent to incoming short-wavelength solar radiation but absorb outgoing long-wavelength infrared radiation from the warm surface. The absorbed radiation is re-emitted in all directions, including back toward the surface, raising the equilibrium surface temperature. As derived in the radiation section, Earth's equilibrium temperature without an atmosphere, after accounting for the reflection of sunlight by clouds and ice (the albedo), would be about 255K (−18°C). The observed mean surface temperature is about 288K (15°C). The 33K greenhouse lift is what makes the planet habitable: without it, the entire surface would be permanently below the freezing point of water.

The important word in this description is equilibrium. The greenhouse effect does not cause an unlimited rise in temperature. It shifts the equilibrium upward: the surface warms until the rate of infrared emission from the top of the atmosphere equals the rate of incoming solar absorption. At that new, higher equilibrium, the planet is in radiative balance again. Increasing the concentration of greenhouse gases shifts the equilibrium to a higher temperature, but the system always finds a new balance. The albedo adds a second layer of complexity: changes in cloud cover or ice extent alter how much sunlight is reflected, which in turn shifts the equilibrium temperature. The interplay between greenhouse warming and albedo reflection is what determines our planet's actual climate.

6. The solar water heater

A solar water heater uses a flat absorber panel, typically painted black (high absorptivity, high emissivity in the infrared), mounted on a roof. Sunlight heats the panel, which in turn heats water flowing through pipes bonded to its surface. The heated water rises naturally to a storage tank mounted above the panel (natural convection, thermosiphon effect), and cooler water from the bottom of the tank flows down to the panel to replace it. No pump is needed: the circulation is driven entirely by the density difference between the hot water rising from the panel and the cooler water descending from the tank.

Three mechanisms work together: radiation from the Sun heats the panel; conduction through the panel material transfers heat to the water; and natural convection circulates the water between the panel and the tank. The system is a self-sustaining heat engine that operates whenever the Sun is shining, converting solar radiation into stored thermal energy with no moving parts and no external power.

7. Why tile feels colder than wood

Ceramic tile and wooden floorboards in the same room are at the same temperature. Yet tile feels cold to bare feet while wood feels warm. The explanation is conduction: tile has a much higher thermal conductivity than wood (roughly $k \approx 1.0 \mathrm{Wm^{-1}K^{-1}}$ for tile versus $k \approx 0.15 \mathrm{Wm^{-1}K^{-1}}$ for wood). So the tile conducts heat away from the foot about seven times faster than wood does. The nerves in the foot respond to the rate of heat loss, not to the floor temperature. A high rate of heat loss triggers a strong "cold" sensation, even though the floor itself is at room temperature.

This is a pure conduction effect. The temperature is the same; the thermal conductivity is different. The lesson: the human body does not measure temperature; it measures heat flux.

Physics earns its keep when it explains the world around you. Let us put these ideas to work.

REAL Example 33

(a) Explain why birds swell their feathers in cold weather. State which mechanism of heat transfer is primarily reduced and what physical change achieves the reduction. (b) Explain why many animals curl into a ball when cold. State which mechanisms of heat transfer are reduced and what physical change achieves the reduction.

Solution

(a) When a bird swells its feathers, the feathers stand away from the body, trapping a thicker layer of air between the feathers and the skin. Air has very low thermal conductivity, so a thicker air layer increases the thermal resistance between the bird's warm body and the cold surroundings. The mechanism primarily reduced is conduction through the insulating air layer. The trapped air is also broken into small pockets that are too narrow for convection currents to develop, so convection through the feather layer is suppressed as well. The physical change is an increase in the effective insulation thickness.

(b) When an animal curls into a ball, it reduces the surface area exposed to the cold surroundings. A sphere has the smallest surface-area-to-volume ratio of any shape. The mechanisms reduced are convection (less surface area for warm air to rise from) and radiation (less surface area to emit radiation, since radiated power is proportional to the emitting area). Curling also reduces exposure to wind, suppressing forced convection. The physical change is a reduction in the effective surface area.

Making Sense of the Answer: *Both behaviours achieve the same goal (reducing heat loss) by different physical means: the bird increases insulation thickness (raises thermal resistance); the animal decreases emitting area (lowers the surface factor in both Stefan-Boltzmann and Newton's law). Neither behaviour generates heat; both work by reducing the rate at which existing body heat escapes.*

Think Like a Physicist: *When answering an application question, always name three things: (1) which mechanism of heat transfer is affected (conduction, convection, radiation, or a combination), (2) what physical quantity in the relevant equation is changed (thickness, area, conductivity, emissivity), and (3) in which direction the change goes (increased insulation, decreased area, lowered emissivity). This three-part structure maps directly to the marking scheme and guarantees full marks.*

HOT Example 34

A room is heated by a domestic iron radiator through which hot water flows at a rate of 0.12kgs^{-1}. The temperature drop of the water between the inlet and the outlet of the radiator is 6.0K. The radiator has a total surface area of 1.5m^2 and its iron walls have a thickness of 2.0mm. The thermal conductivity of iron is $k = 80\text{Wm}^{-1}\text{K}^{-1}$ and the specific heat capacity of water is $c = 4200\text{Jkg}^{-1}\text{K}^{-1}$. (a) Calculate the rate at which heat is supplied to the room. (b) Calculate the mean temperature difference across the iron walls of the radiator.

Solution

(a) The rate of heat supplied equals the rate at which the water loses thermal energy as it flows through the radiator:

$$\frac{dQ}{dt} = \frac{m}{t} c\,\Delta\theta$$

$$= 0.12\text{kgs}^{-1} \times 4200\text{Jkg}^{-1}\text{K}^{-1} \times 6.0\text{K} = 3024\text{W} \approx 3.0\text{kW}$$

The radiator supplies about 3.0kW of heat to the room.

(b) This heat must conduct through the iron walls of the radiator to reach the room air. By Fourier's law:

$$\frac{dQ}{dt} = \frac{kA\Delta\theta}{x}$$

Rearranging for the temperature difference across the walls:

$$\Delta\theta = \frac{\frac{dQ}{dt} \times x}{kA} = \frac{3024\text{W} \times 2.0 \times 10^{-3}\,\text{m}}{80\text{Wm}^{-1}\text{K}^{-1} \times 1.5\text{m}^2} = 0.050\text{K}$$

The mean temperature difference across the iron walls is only about 0.050K, or 50 millikelvin.

Making Sense of the Answer: *A temperature difference of just 0.05K across the iron walls confirms what we observed in the radial-conduction section: metal walls are essentially transparent to heat. The thermal resistance of 2mm of iron is so tiny that it barely impedes the 3kW heat flow. The real bottleneck for heating the room is the convective heat transfer from the outer surface of the radiator to the room air, not the conduction through the metal. This is why radiator design focuses on maximising surface area (more fins) rather than reducing wall thickness.*

Think Like a Physicist: *This problem combines two different heat-transfer equations (water calorimetry for part (a), Fourier conduction for part (b)) on the same device. The key insight is that the heat-flow rate is the same in both equations: what the water loses is what the walls must conduct. Linking two equations through a shared quantity is a common exam technique, and it always starts with identifying what quantity is the same in both.*

This chapter has taken the three heat-transfer mechanisms from their definitions to their mathematical laws, from their laboratory measurements to their everyday applications. Conduction, convection, and radiation are not three separate subjects; they are three faces of the same physical question: how does thermal energy move from one place to another? In every real situation, all three act together, and the physicist's job is to identify which one dominates and why. The tools built in this

chapter, from Fourier's law to Stefan-Boltzmann's law to Newton's law of cooling, are the instruments for answering that question in any context the examination or the world can present. The miscellaneous worked examples that follow put all of these instruments on the table at once.

MISCELLANEOUS WORKED EXAMPLES ON HEAT TRANSFER

Example 35

(a) A brick wall separating a warm room from the cold outdoors feels cold to the touch even though the room air is at 22°C. Explain which heat-transfer mechanism is responsible for the wall surface being below room temperature and why the wall surface cannot be at the same temperature as the room air.

(b) A brick wall of thickness 20cm, thermal conductivity $k = 0.72\text{Wm}^{-1}\text{K}^{-1}$, and area 10m² separates a room at 22°C from outside air at 5°C. The inner surface has a convective film resistance of $0.12\text{m}^2\text{KW}^{-1}$ and the outer surface has a convective film resistance of $0.06\text{m}^2\text{KW}^{-1}$. (i) Calculate the total thermal resistance of the wall including both surface films. (ii) Calculate the rate of heat loss through the wall.

Solution

(a) The room air transfers heat to the wall surface by convection. This convective boundary layer has a finite thermal resistance (called the surface film resistance), meaning there must be a temperature difference between the room air and the wall surface to drive the heat flow. The wall surface is therefore always slightly cooler than the room air. The larger the rate of heat flow through the wall, the larger this temperature drop across the film. The wall cannot be at room temperature because that would imply zero temperature difference across the film, which would mean zero heat flow, contradicting the fact that heat is continuously flowing through the wall to the colder outdoors.

(b) Three thermal resistances in series: inner film, brick wall, outer film.

Brick wall insulance: $\frac{x}{k} = \frac{0.20\text{m}}{0.72\text{Wm}^{-1}\text{K}^{-1}} = 0.278\text{m}^2\text{KW}^{-1}$

Total insulance in series: $(0.12 + 0.278 + 0.06)\text{m}^2\text{KW}^{-1} = 0.458\text{m}^2\text{KW}^{-1}$

Total resistance for area $A = 10\text{m}^2$:

$$R = \frac{r}{A} = \frac{0.458\text{m}^2\text{KW}^{-1}}{10\text{m}^2} = 0.0458\text{KW}^{-1}$$

(ii) Rate of heat loss:

$$\frac{dQ}{dt} = \frac{\Delta T}{R} = \frac{22°\text{C} - 5°\text{C}}{0.0458\text{KW}^{-1}} = \frac{17\text{K}}{0.0458\text{KW}^{-1}} = 371\text{W}$$

Example 36

(a) A student claims a vacuum flask *"keeps things hot."* Another student says it *"keeps things cold."* Explain how both statements are correct using the physics of the three heat-transfer mechanisms.

(b) Tea at 90°C in a thermos flask is placed in a room at 20°C. After 3 hours the tea has cooled to 85°C. (i) Find the cooling constant λ. (ii) How long does it take for the tea to reach 50°C? (iii) What is the half-cooling time?

Solution

(a) The vacuum flask suppresses all three heat-transfer mechanisms: the vacuum eliminates conduction and convection across the gap, the silvered surfaces reduce radiation, and the stopper prevents convection at the opening. These mechanisms operate identically regardless of the direction of heat flow. When a hot liquid is inside, the flask slows heat escaping outward. When a cold liquid is inside, the same barriers slow heat entering from the warm surroundings. The flask does not know whether it is keeping heat in or keeping heat out; it simply reduces the rate of heat transfer in both directions by the same set of mechanisms.

(b)(i) Using the logarithmic form of Newton's law:

$$\lambda = \frac{\ln\left(\frac{90°\text{C} - 20°\text{C}}{85°\text{C} - 20°\text{C}}\right)}{3\text{hr}} = 0.0247\text{hr}^{-1}$$

(ii) Time to reach 50°C:

$$t = \frac{\ln\left(\frac{90°\text{C} - 20°\text{C}}{50°\text{C} - 20°\text{C}}\right)}{\lambda} = \frac{\ln\left(\frac{70\text{K}}{30\text{K}}\right)}{0.0247\text{hr}^{-1}} = 34.3\text{hr}$$

(iii) Half-cooling time:

$$t_{1/2} = \frac{\ln 2}{\lambda} = \frac{0.693}{0.0247\text{hr}^{-1}} = 28.1\text{hr}$$

Example 37

(a) A blacksmith heats a piece of iron continuously from room temperature. Describe and explain the colour changes the blacksmith observes, linking each change to the behaviour of the spectral emission curve.

(b) A tungsten filament operates at 2800K. (i) Find the peak emission wavelength. (ii) The visible spectrum spans approximately 400nm to 700nm. Is the peak within this range? What does this tell you about the efficiency of incandescent bulbs? (iii) If the filament temperature is increased to 3500K, by what factor does the total radiated power per unit area increase?

Solution

(a) At room temperature the iron emits only in the infrared and appears dark. As it heats, the peak of its spectral emission curve shifts to shorter wavelengths (Wien's law). At about 800K the peak enters the far-red edge of the visible spectrum and the iron glows dull red. Further heating shifts the peak through the visible range: bright red, orange, then yellow. At each stage the peak also grows taller (as the fifth power of temperature), so the iron becomes visibly brighter. At very high temperatures the peak moves into the blue region, but because the curve is broad, significant radiation is emitted at all visible wavelengths simultaneously, and the iron appears white rather than blue.

(b)(i) Wien's law:

$$\lambda_{max} = \frac{b}{T} = \frac{2.90 \times 10^{-3}\,\text{mK}}{2800\text{K}} = 1036\text{nm}$$

(ii) The peak is at 1036nm, which is in the near-infrared, well outside the visible range (400-700nm). Most of the filament's radiated energy comes out as invisible heat rather than visible light. This makes incandescent bulbs very inefficient as light sources: only a small fraction of the total radiated power falls within the visible spectrum.

(iii) Stefan-Boltzmann power ratio:

$$\frac{P_2}{P_1} = \left(\frac{3500\text{K}}{2800\text{K}}\right)^4 = (1.25)^4 = 2.44$$

The total radiated power per unit area increases by a factor of about 2.4.

Example 38

(a) A student performing Searle's experiment removes the lagging from the conducting bar to "see what happens." Explain how this error affects the calculated value of thermal conductivity and in which direction (too high or too low).

(b) In a Searle's apparatus experiment on a brass bar, the following data are recorded: distance between thermometers $x = 25\text{cm}$, bar diameter $d = 3.0\text{cm}$, bar thermometer readings $\theta_1 = 97°\text{C}$ and $\theta_2 = 84°\text{C}$, water flow rate $\frac{m}{t} = 2.4\text{gmin}^{-1}$, water inlet temperature $\theta_3 = 18°\text{C}$, water outlet temperature $\theta_4 = 31°\text{C}$. Calculate the thermal conductivity of brass.

Solution

(a) Without lagging, heat escapes from the sides of the bar into the surrounding air. Less heat reaches the water coil at the cold end than actually entered the bar at the hot end. The water warms less than it should, so the measured heat-flow rate (from the water temperature rise) is lower than the true conduction rate through the bar. Since the calculation assumes all conducted heat reaches the water, the equation underestimates the true heat flow, and the calculated value of k comes out too low.

(b) Converting units:

$x = 0.25\text{m}$, $A = \pi(0.015\text{m})^2 = 7.07 \times 10^{-4}\,\text{m}^2$

$$\frac{m}{t} = 2.4\text{gmin}^{-1} = 4.0 \times 10^{-5}\,\text{kgs}^{-1}$$

From Searle's working equation:

$$k = \frac{\frac{m}{t}\,c(\theta_4 - \theta_3)\,x}{A(\theta_1 - \theta_2)}$$

$$= \frac{4.0 \times 10^{-5}\,\text{kgs}^{-1} \times 4200\text{Jkg}^{-1}\text{K}^{-1} \times 13\text{K} \times 0.25\text{m}}{7.07 \times 10^{-4}\,\text{m}^2 \times 13\text{K}} = 59\text{Wm}^{-1}\text{K}^{-1}$$

The thermal conductivity of brass is approximately $59\text{Wm}^{-1}\text{K}^{-1}$.

Example 39

(a) In the Lee's disc method, the brass slab Y must be cooled alone after the steady state has been reached. Explain why the cooling curve cannot be obtained while the steam chest is still in contact with the sample.

(b) A Lee's disc experiment on a cardboard disc gives: disc thickness $x = 3.2\text{mm}$, disc diameter $d = 10.0\text{cm}$, steam chest temperature $\theta_1 = 99°\text{C}$, brass slab temperature $\theta_2 = 54°\text{C}$, mass of brass slab $m = 0.80\text{kg}$, specific heat of brass $c = 380\text{Jkg}^{-1}\text{K}^{-1}$, gradient of cooling curve at θ_2 is 0.32Kmin^{-1}. Calculate the thermal conductivity of cardboard.

Solution

(a) While the steam chest is in contact, heat is continuously flowing into slab Y through the sample, maintaining Y at its steady-state temperature. The slab is not cooling; it is in thermal equilibrium between the heat arriving from the sample and the heat it loses to the air. To measure the rate at which Y loses heat at temperature θ_2, we need Y to be cooling freely with no heat input. Removing the steam chest and sample ensures that the only process occurring is heat loss from Y to the surroundings, and the gradient of the resulting cooling curve at θ_2 gives exactly the rate of heat loss that was, during the steady state, being supplied by conduction through the sample.

(b) Converting units:

$$x = 3.2 \times 10^{-3}\,\text{m}, A = \pi(0.050\text{m})^2 = 7.854 \times 10^{-3}\,\text{m}^2$$

$$\frac{a}{b} = 0.32\text{Kmin}^{-1} = 5.33 \times 10^{-3}\,\text{Ks}^{-1}$$

From Lee's disc working equation:

$$k = \frac{mc\frac{a}{b}x}{A(\theta_1 - \theta_2)} = \frac{0.80\text{kg} \times 380\text{Jkg}^{-1}\text{K}^{-1} \times 5.33 \times 10^{-3}\,\text{Ks}^{-1} \times 3.2 \times 10^{-3}\,\text{m}}{7.854 \times 10^{-3}\,\text{m}^2 \times 45\text{K}} = 0.015\text{Wm}^{-1}\text{K}^{-1}$$

The thermal conductivity of cardboard is approximately $0.015\text{Wm}^{-1}\text{K}^{-1}$.

Example 40

(a) A nurse argues: *"The patient loses heat by radiation only to objects colder than themselves."* Evaluate this claim using Prevost's theory of heat exchange.

(b) A metal body of surface area 0.050m^2 and emissivity 0.60 sits in surroundings at 300K. (i) Calculate the net rate of radiative heat loss when the body is at 600K, at 400K, and at 310K. (ii) At which of these temperatures does the surroundings' contribution to the absorbed power exceed 10% of the body's emitted power? (iii) Below what excess temperature would Newton's law be a reasonable approximation?

Solution

(a) The claim is wrong. Prevost's theory states that every body emits and absorbs radiation continuously and simultaneously, regardless of the temperatures of the surrounding objects. The patient emits radiation toward all objects in the room, hot or cold, and simultaneously absorbs radiation from all of them. The net heat loss is the difference between emitted and absorbed power. A patient in a room full of objects at the same temperature still emits radiation; the reason there is no net loss at equilibrium is that the patient absorbs exactly as much as they emit, not that they stop emitting. The nurse's error is confusing net transfer with emission alone.

(b)(i) Using the net Stefan-Boltzmann formula at each temperature:

At 600K: $P_{net} = 0.60 \times 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4} \times 0.050\text{m}^2 \times ((600\text{K})^4 - (300\text{K})^4) = 207\text{W}$

At 400K: $P_{net} = 1.701 \times 10^{-9} \times (2.56 \times 10^{10} - 8.1 \times 10^9)\,\text{W} = 29.7\text{W}$

At 310K: $P_{net} = 1.701 \times 10^{-9} \times (9.24 \times 10^9 - 8.1 \times 10^9)\,\text{W} = 1.94\text{W}$

(ii) At 600K, the emitted power is $1.701 \times 10^{-9} \times 1.296 \times 10^{11} = 220\text{W}$ and the absorbed power is $1.701 \times 10^{-9} \times 8.1 \times 10^9 = 13.8\text{W}$. The ratio is 13.8/220 = 6.3%, below 10%. At 400K, emitted = 43.5W, absorbed = 13.8W, ratio = 31.7%, well above 10%. At 310K, emitted = 15.7W, absorbed = 13.8W, ratio = 87.8%. The surroundings term exceeds 10% at 400K and 310K.

(iii) Newton's law approximation requires the excess to be small compared to the absolute surroundings temperature. With surroundings at 300K, the threshold is roughly 10% of 300K = 30K. So Newton's law is reasonable for body temperatures below about 330K (excess below about 30K).

Example 41

(a) In rural Tanzania, water stored in a porous clay pot outdoors can become noticeably cooler than the surrounding air temperature, while water stored in a sealed plastic jerry can stays at air temperature. Explain why the clay pot cools the water below the air temperature, identifying the mechanism responsible.

(b) A blackbody metal sphere of surface area 0.030m^2, mass 2.0kg, and specific heat capacity $c = 400\text{Jkg}^{-1}\text{K}^{-1}$ is at 350K in surroundings at 300K. (i) Calculate the net rate of radiative heat loss. (ii) Hence calculate the initial rate of fall of temperature of the sphere.

Solution

(a) The clay pot is porous, so water slowly seeps through the walls to the outer surface and evaporates. Evaporation requires latent heat, which is drawn from the water and the pot itself, cooling them. This heat extraction continues as long as the outer surface remains moist and the surrounding air is not saturated with water vapour. The water temperature can fall below the air temperature because the evaporative cooling removes heat faster than convection and radiation from the warm air can replace it. The system reaches a dynamic equilibrium where the rate of evaporative cooling equals the rate of heat gain from

the surroundings, at a temperature below the air temperature. A sealed plastic jerry can has no evaporation from its surface, so the only heat-transfer mechanisms are conduction, convection, and radiation between the water and the surroundings, which drive the water toward air temperature, not below it.

(b)(i) Net radiative heat loss:

$$P_{net} = \varepsilon\sigma A\left(T^4 - T_s{}^4\right) = 1 \times 5.67 \times 10^{-8}\,Wm^{-2}K^{-4} \times 0.030m^2 \times ((350K)^4 - (300K)^4) = 11.7W$$

(ii) Rate of temperature fall:

From $P_{net} = mc\frac{d\theta}{dt}$;

$$\frac{d\theta}{dt} = \frac{P_{net}}{mc} = \frac{11.7W}{2.0kg \times 400Jkg^{-1}K^{-1}} = 0.0146Ks^{-1}$$

Example 42

(a) Two identical cups of tea start at 80°C. One is placed in a room at 10°C, the other in a room at 20°C. The cups have the same mass, surface area, and emissivity. Do they have the same half-cooling time? Explain your reasoning.

(b) A blackbody sphere of radius 5.0cm, mass 4.1kg, and specific heat capacity $c = 450Jkg^{-1}K^{-1}$ sits in a room at 300K. (i) Find the cooling constant λ for this sphere. (ii) Hence find the half-cooling time.

Solution

(a) No, they do not have the same half-cooling time. The Newton cooling constant k depends on the cube of the surroundings temperature: $k = 4\varepsilon\sigma AT_s{}^3$. Since the cooling constant $\lambda = k/(mc)$ depends on $T_s{}^3$, the cup in the 10°C room (283K) has a smaller λ than the cup in the 20°C room (293K). The half-cooling time $t_{1/2} = \ln 2 / \lambda$ is therefore longer in the colder room. This may seem counterintuitive (a colder room should cool things faster), but it reflects the radiation physics: the room at 20°C radiates more heat back to the cup, but the cup also emits more net power because the Stefan-Boltzmann net formula involves $T_s{}^3$ in the linearised Newton's constant. The dominant effect is that the radiative coupling between the cup and the surroundings is stronger at higher surroundings temperatures, giving a larger λ and therefore a shorter half-cooling time.

(b)(i) Newton's law of cooling is a small-excess approximation to Stefan-Boltzmann. The **heat**-loss constant for a body cooling by radiation is:

$$k = 4\varepsilon\sigma AT_s{}^3$$

Surface area of the sphere:

$$A = 4\pi(0.050m)^2 = 0.0314m^2$$

Substituting:

$$k = 4 \times 1 \times 5.67 \times 10^{-8}\,Wm^{-2}K^{-4} \times 0.0314m^2 \times (300K)^3 = 0.192WK^{-1}$$

The cooling constant:

$$\lambda = \frac{k}{mc} = \frac{0.192WK^{-1}}{4.1kg \times 450Jkg^{-1}K^{-1}} = 1.04 \times 10^{-4}\,s^{-1}$$

(ii) The half-cooling time:

$$t_{1/2} = \frac{\ln 2}{\lambda} = \frac{0.693}{1.04 \times 10^{-4}\,s^{-1}} = 6660s = 111min \approx 1.85hr$$

Example 43

(a) A body at 90°C cools freely in still air at 20°C. Explain why the rate of temperature fall is not constant even though the surroundings temperature remains at 20°C throughout.

(b) A block of ice of mass 200g at 0°C is placed in a room at 25°C. The rate of heat absorption obeys Newton's law of cooling with $k = 1.2WK^{-1}$. The specific latent heat of fusion of ice is $L_f = 334kJkg^{-1}$. (i) Calculate the rate of heat absorption by the ice. (ii) Calculate the time taken for the ice to melt completely. (iii) Explain why the rate of heat absorption remains constant during melting.

Solution

(a) Newton's law of cooling says the rate of temperature fall is proportional to the excess temperature over the surroundings. At the start, the excess is 70°C, so the rate of cooling is large. As the body cools, the excess shrinks, so the driving force for cooling weakens and the rate of temperature fall decreases. Each degree of cooling removes some of the driving force that caused the cooling in the first place. This self-slowing is the hallmark of exponential decay: the rate depends on the current state, not on the initial state.

(b)(i) The surroundings (25°C) are hotter than the ice (0°C), so heat flows from the surroundings to the ice. The excess of the surroundings over the ice is:

$$\frac{dQ}{dt} = k(\theta_s - \theta) = 1.2WK^{-1} \times (25°C - 0°C) = 1.2WK^{-1} \times 25K = 30W$$

(ii) Heat needed to melt the ice:

$$Q = mL_f = 0.200\text{kg} \times 334 \times 10^3\ \text{Jkg}^{-1} = 66800\text{J}$$

$$t = \frac{Q}{\frac{dQ}{dt}} = \frac{66800\text{J}}{30\text{W}} = 2227\text{s} = 37.1\text{min}$$

(iii) During melting, the ice stays at 0°C because the absorbed heat goes into breaking intermolecular forces (latent heat) rather than raising the temperature. Since the ice temperature does not change, the excess temperature remains constant at 25°C, and Newton's law of cooling gives a constant rate of heat absorption. This is unlike normal cooling, where the temperature changes and the rate decays exponentially. During a phase change, Newton's law produces a constant heat flow, not an exponential one.

Example 44

(a) Provide a mathematical explanation for why the body's temperature θ never actually reaches the surrounding temperature θ_s on the Newton cooling curve.

(b) A double-glazed window consists of two glass panes, each 4.0mm thick ($k_{glass} = 0.80\text{Wm}^{-1}\text{K}^{-1}$), separated by a 10mm air gap ($k_{air} = 0.025\text{Wm}^{-1}\text{K}^{-1}$). The window area is 1.5m². Inside temperature 22°C, outside temperature 5°C. (i) Calculate the thermal resistance of each of the three layers. (ii) Calculate the total thermal resistance and the rate of heat loss. (iii) Calculate the rate of heat loss through a single glass pane of thickness 8.0mm (same area) and compare.

Solution

(a) The integrated form of Newton's law is $\theta = \theta_s + (\theta_i - \theta_s)e^{-\lambda t}$. For the body to reach θ_s, we would need $(\theta_i - \theta_s)e^{-\lambda t} = 0$. Since $(\theta_i - \theta_s)$ is a positive constant and $e^{-\lambda t}$ is always positive for any finite t, their product can never be zero. The exponential function decays toward zero but never reaches it. Therefore the body approaches θ_s ever more closely but never arrives: the surroundings temperature is an asymptote, not a destination.

(b)(i) Thermal resistance of each layer (R = x/(kA)):

Each glass pane: $R_{glass} = \frac{4.0 \times 10^{-3}\ \text{m}}{0.80\text{Wm}^{-1}\text{K}^{-1} \times 1.5\text{m}^2} = 3.33 \times 10^{-3}\ \text{KW}^{-1}$

Air gap: $R_{air} = \frac{10 \times 10^{-3}\ \text{m}}{0.025\text{Wm}^{-1}\text{K}^{-1} \times 1.5\text{m}^2} = 0.267\text{KW}^{-1}$

(ii) Total resistance and heat loss:

$$R_{total} = (3.33 \times 10^{-3} + 0.267 + 3.33 \times 10^{-3})\text{KW}^{-1} = 0.274\text{KW}^{-1}$$

$$\frac{dQ}{dt} = \frac{\Delta T}{R} = \frac{(22-5)\text{K}}{0.274\text{KW}^{-1}} = 62\text{W}$$

(iii) Single 8mm glass pane:

$$R_{single} = \frac{8.0 \times 10^{-3}\ \text{m}}{0.80\text{Wm}^{-1}\text{K}^{-1} \times 1.5\text{m}^2} = 6.67 \times 10^{-3}\ \text{KW}^{-1}$$

$$\frac{dQ}{dt} = \frac{\Delta T}{R} = \frac{(22-5)\text{K}}{6.67 \times 10^{-3}\ \text{KW}^{-1}} = 2550\text{W}$$

Have you noticed? The double-glazed window loses 62W; the single pane loses 2550W. The double-glazing reduces heat loss by a factor of about 41. The air gap, which is only 10mm thick, provides 97.6% of the total thermal resistance.

Example 45

(a) A solar water heater panel is painted black rather than white. Explain why, discussing what happens to both the absorptivity and the emissivity of the surface, and why the increased emissivity is an acceptable trade-off.

(b) A flat-plate solar water heater has an absorber area of 2.0m² with absorptivity 0.90. The solar intensity on a clear day is 800Wm^{-2}. The tank loses heat to the surroundings (air at 25°C) by Newton's law with $k = 8.0\text{WK}^{-1}$. (i) Calculate the power absorbed by the panel. (ii) Find the equilibrium temperature of the water. (iii) On a cloudy day the solar intensity drops to 300Wm^{-2}. Find the new equilibrium temperature.

Solution

(a) A black surface has high absorptivity (close to 1), meaning it absorbs most of the incoming solar radiation and converts it to thermal energy. A white surface reflects most of the sunlight and absorbs very little. Painting the panel black maximises the solar energy captured. However, by Kirchhoff's law, a good absorber is also a good emitter: the black panel will also radiate more infrared energy than a white one would. This increased emissivity means the panel loses more heat by radiation at high temperatures. The trade-off is acceptable because the dominant heat loss from a solar panel at typical operating temperatures (40-70°C) is convection to the surrounding air, not radiation. The gain from absorbing more sunlight far outweighs the modest increase in radiative loss.

(b)(i) Power absorbed:

$$P_{abs} = 0.90 \times 800\text{Wm}^{-2} \times 2.0\text{m}^2 = 1440\text{W}$$

(ii) At equilibrium, power absorbed = power lost:

$$P_{abs} = k(\theta - \theta_s)$$

$$1440\text{W} = 8.0\text{WK}^{-1} \times (\theta - 25°\text{C})$$
$$\theta - 25°\text{C} = \frac{1440\text{W}}{8.0\text{WK}^{-1}} = 180\text{K}$$
$$\theta = 25°\text{C} + 180°\text{C} = 205°\text{C}$$

(iii) Cloudy day:

$$P_{abs} = 0.90 \times 300\text{Wm}^{-2} \times 2.0\text{m}^2 = 540\text{W}$$
$$\theta = 25°\text{C} + \frac{540\text{W}}{8.0\text{WK}^{-1}} = 25°\text{C} + 67.5°\text{C} = 92.5°\text{C}$$

Eleven examples have stretched every tool built in this chapter across problems that no single section could have posed. The Digging Deeper questions that follow push further still.

DIGGING DEEPER EXERCISE 3

EXERCISE 3A: BINDER QUESTIONS

Question 1

A metal spoon and a wooden *mwiko* are both resting in a pot of boiling water. A cook grabs the metal spoon handle for half a second and burns her fingers, then grabs the wooden mwiko handle and feels no pain.

(a) Both handles were in contact with the same boiling water. Explain why the metal handle is painful even for a brief contact while the wooden handle is comfortable.
(b) A student says: *"The metal spoon is hotter than the wooden mwiko."* Is this correct? Explain.

Question 2

Air has a lower thermal conductivity ($k \approx 0.026 \text{Wm}^{-1}\text{K}^{-1}$) than felt ($k \approx 0.04 \text{Wm}^{-1}\text{K}^{-1}$). Yet wrapping a hot pipe in felt reduces heat loss far more effectively than leaving an open air gap of the same thickness.

(a) Explain why a material with higher thermal conductivity can be a better insulator than one with lower thermal conductivity.
(b) Identify the heat-transfer mechanism that the felt suppresses but the open air gap does not.
(c) Explain why fibre-based insulation materials (felt, glass wool, feathers) all work by the same physical principle.

Question 3

A wall consists of three layers of different materials in series. In the steady state, the temperature drop across each layer is different.

(a) Using the thermal Ohm's law analogy, explain why the temperature drop across each layer is proportional to that layer's thermal resistance.
(b) Which layer has the largest temperature drop: the one with the highest thermal conductivity, the lowest thermal conductivity, or the greatest thickness? Explain.
(c) A builder argues: *"The thickest layer always provides the most insulation."* Explain why this claim is not necessarily correct.

Question 4

A copper bar has one end in steam at 100°C and the other end in melting ice at 0°C.

(a) If the bar is perfectly lagged along its entire length, sketch the temperature profile along the bar and explain why it has this shape.
(b) The lagging is now removed. Sketch the new temperature profile on the same axes. Explain why the profile curves and in which direction it curves.
(c) In the unlagged case, is the temperature at the midpoint of the bar above or below 50°C? Explain.

Question 5

Newton's law of cooling applies under forced convection at any excess temperature, or under natural convection when the excess temperature is below about 30°C.

(a) A beaker of water at 65°C cools in a room at 28°C under natural convection. Identify the temperature at which the cooling enters the Newton regime.
(b) During the early part of the cooling (when the excess exceeds 30°C), does the actual cooling rate become faster or slower than Newton's law would predict? Explain.

Question 6

A body at 1200K loses heat by radiation in surroundings at 300K. A student says: *"If I cool the body to 600K, the net radiative heat loss will be exactly one-sixteenth of the original, because Stefan-Boltzmann's law involves T^4."* Explain the error in the student's reasoning.

Question 7

A vacuum flask minimises heat transfer by all three mechanisms.

(a) A small hole is accidentally made in the stopper. Which mechanism of heat transfer does this most immediately increase? Explain.
(b) Over many years the silvered coating on the inner walls deteriorates. Explain how this affects the rate of heat transfer and which mechanism is involved.
(c) A student argues: *"This flask keeps my coffee hot, so it cannot keep cold juice cold."* Explain the error in this reasoning.

Question 8

Two identical metal blocks are placed in the same room at 20°C. Block A starts at 80°C and block B starts at 60°C.

(a) Which block cools faster initially? Explain.
(b) Which block takes longer to come within 2°C of room temperature? Explain.
(c) Do the two blocks have the same half-cooling time? Explain whether your answers to (a) and (b) are contradictory.

Question 9

A premature baby in a hospital incubator can lose dangerous amounts of heat by radiation even when the air inside the incubator is at the correct temperature of 36°C.

(a) Using Prevost's theory, explain how the baby can experience a net radiative heat loss even when the air temperature is correct.
(b) Identify the specific physical feature of the incubator that causes this problem. What temperature would the walls need to be at to eliminate the net radiative loss?
(c) Modern incubators use double-walled enclosures with warm water circulating between the walls. Explain how this design solves the radiation problem.

Question 10

On a cold night, a layer of ice forms on a pond. As the ice thickens over many hours, its rate of growth becomes progressively slower, even though the air temperature above and the water temperature below remain constant.

(a) Identify the heat-transfer mechanism that carries heat from the water below through the ice to the cold air above.
(b) Explain why the rate of heat flow through the ice decreases as the ice becomes thicker, even though the temperature difference across it stays the same.
(c) Explain why the ice therefore acts as a self-limiting insulator: the thicker it grows, the harder it becomes for it to grow further.

Question 11

A body is cooling in a room at 25°C according to Newton's law. After some time, the body (still well above room temperature) is suddenly moved to a second room at 10°C.

(a) Describe what happens to the rate of temperature fall at the instant of the transfer. Does it increase, decrease, or stay the same? Explain.
(b) Sketch the complete cooling curve from the start in the 25°C room through the transfer and into the 10°C room, marking the transfer point and both asymptotes clearly.
(c) Does the body have the same cooling constant λ in both rooms? Explain.

Question 12

The International Space Station orbits Earth in the near-vacuum of space. Its electronic equipment generates significant heat that must be removed to prevent overheating.

(a) Explain why natural convection cannot cool the equipment inside the station, even though the cabin is filled with air at normal atmospheric pressure.
(b) Fans circulate air over hot components. What type of convection is this, and why does it work in microgravity when natural convection does not?
(c) On the outside of the station, large radiator panels radiate waste heat into space. Explain why conduction and convection are both useless for removing heat from the exterior, and why radiation is the only viable mechanism.

EXERCISE 3B: REAL QUESTIONS

Question 13

Kipanga visits his aunt's new house in Dodoma. Walking barefoot from the wooden bedroom floor onto the marble bathroom floor, he exclaims: *"The bathroom is freezing! Someone left the window open!"* His aunt laughs: *"The window is shut. Both floors are at the same temperature."*

(a) Explain to Kipanga why the marble floor feels colder than the wooden floor even though a thermometer placed on each surface would read the same temperature.
(b) Kipanga puts on rubber sandals and steps onto the marble again. He says it feels much less cold. Explain why, identifying the physical property of rubber that is responsible.

Question 14

A cook in a Tanzanian kitchen heats a clay pot on a charcoal *jiko*. The pot has a wooden handle.

(a) Explain why the clay body of the pot becomes extremely hot while the wooden handle stays cool enough to grip, even though both are attached to the same pot.
(b) The cook wraps a damp cloth around the handle as an extra precaution. Using two different heat-transfer mechanisms, explain why the damp cloth keeps her hand cooler than a dry cloth would.

Question 15

A blanket keeps a sleeping person warm on a cold night. The same blanket, wrapped around a block of ice, slows the melting of the ice on a hot day.

(a) Does the blanket generate any heat? Explain.
(b) Identify all three heat-transfer mechanisms that the blanket reduces and explain how it reduces each.
(c) Explain why the blanket works in both directions (keeping heat in and keeping heat out) using a single physical principle.

Question 16

A solar water heater installed on a rooftop in Mwanza uses a flat black absorber panel connected to a storage tank mounted above it. There is no electric pump.

(a) Explain why the absorber panel is painted black rather than left as shiny aluminium.
(b) Water circulates automatically between the panel and the tank without any pump. Identify the mechanism that drives this circulation and explain how it works.
(c) The heater is more effective on a clear sunny day than on an overcast day at the same air temperature. Explain why.

Question 17

A student writes in an essay: *"The greenhouse effect permanently traps heat inside Earth's atmosphere, preventing any heat from escaping to space. This means Earth's temperature will rise indefinitely."*

(a) Identify the two errors in the student's reasoning.
(b) Explain the correct equilibrium mechanism that determines Earth's surface temperature.
(c) Predict what would happen to Earth's surface temperature if all greenhouse gases were suddenly removed from the atmosphere.

Question 18

On a cold night in Arusha, a cup of hot tea at 80°C and a glass of cold water at 10°C are left on the same table in a room at 20°C.

(a) Which object loses heat faster initially? Explain.
(b) Which object gains heat initially? Explain.
(c) Will both objects eventually reach the same temperature? If so, what temperature, and which one reaches it first? Explain.

Question 19

In Dar es Salaam, daytime temperatures regularly exceed 30°C and houses with corrugated iron roofs can become unbearably hot inside.

(a) Propose one modification to the building that would reduce heat entering by conduction, and explain the physics behind your proposal.
(b) Propose one modification that would reduce heat entering by convection, and explain the physics.
(c) Propose one modification that would reduce heat entering by radiation, and explain the physics.

Question 20

A mother in Iringa puts a pot of *ugali* on a charcoal *jiko* to keep it warm for a latecomer. She then wraps the pot in several layers of thick newspaper.

(a) Identify the two mechanisms of heat transfer that the newspaper wrap primarily reduces. Explain how it reduces each.
(b) After 30 minutes the *ugali* is still hot. After 2 hours it has cooled noticeably. Using Newton's law of cooling, explain why the rate of cooling is fast at first and slows down over time, even though the newspaper stays the same.
(c) Would wrapping the pot in aluminium foil instead of newspaper keep the *ugali* hot for longer? Explain which mechanism each material is better at suppressing.

Question 21

Kipute pours equal amounts of boiling tea into a metal mug and a ceramic mug. After five minutes she touches both mugs. The metal mug feels cooler than the ceramic mug. She says: *"The metal mug cooled the tea faster because metal has higher thermal conductivity."*

(a) Is Kipute correct that the tea in the metal mug cooled faster? Explain.
(b) Kipute's reasoning uses thermal conductivity. Explain why thermal conductivity alone does not fully explain the different cooling rates. What other property of the mug material matters?
(c) If both mugs were perfectly lagged on the outside so that no heat could escape, would the tea in both mugs be at the same temperature after five minutes? Explain.

Question 22

A builder adds a 5cm layer of insulating foam to the inside surface of a concrete roof. Before the foam was added, the occupants complained that the house was hot during the day and cold at night.

(a) Explain how the foam reduces the rate of heat flow through the roof during the day, identifying the physical property of the foam that makes it effective.
(b) The builder argues that the greatest temperature drop in the composite roof-plus-foam system occurs across the foam layer, not across the much thicker concrete. Explain why this is correct.
(c) At night, when the outside is cooler than the inside, the same foam layer now keeps the house warmer. Explain why the insulation works in both directions.

EXERCISE 3C: HOT QUESTIONS

Useful constants: Stefan-Boltzmann constant $\sigma = 5.67 \times 10^{-8}\,\text{Wm}^{-2}\text{K}^{-4}$; Wien's displacement constant $b = 2.90 \times 10^{-3}\,\text{mK}$.

Question 23

A composite lagged rod consists of a copper section (length 0.50m, $k = 380\text{Wm}^{-1}\text{K}^{-1}$) joined to a section of unknown material (length 0.50m). Both sections have cross-sectional area $2.0 \times 10^{-4}\,\text{m}^2$. The hot end of the copper is at 200°C, the cold end of the unknown material is at 20°C, and the interface temperature is 170°C. Find the thermal conductivity of the unknown material.

Question 24

A vertical soil column consists of a sand layer (53cm, $k = 0.35\text{Wm}^{-1}\text{K}^{-1}$) on top of a loam layer (29cm, $k = 0.75\text{Wm}^{-1}\text{K}^{-1}$). The cross-sectional area is 1.0m². The top surface is at 37°C and the bottom at 14°C. Find the heat flux through the column and the temperature at the sand-loam interface. After light rain, the conductivities change to $k_{\text{wet sand}} = 1.8\text{Wm}^{-1}\text{K}^{-1}$ and $k_{\text{wet loam}} = 2.5\text{Wm}^{-1}\text{K}^{-1}$. By what factor does the heat flow increase?

Question 25

A body cools from temperature θ_1 to θ_2 in time t_1, and from θ_2 to θ_3 in time t_2, where the temperature drops are equal: $\theta_1 - \theta_2 = \theta_2 - \theta_3$. The surroundings are at θ_s. Using Newton's law of cooling, show algebraically that $t_2 > t_1$.

Question 26

A small blackbody sphere is heated to 500K and suspended in a vacuum at $T_s = 300\text{K}$. The sphere has mass 0.50kg, specific heat capacity $c = 400\text{Jkg}^{-1}\text{K}^{-1}$, and surface area $A = 3.0 \times 10^{-3}\,\text{m}^2$.

(a) Write the exact differential equation for the rate of temperature change of the sphere, using Stefan-Boltzmann's law directly.
(b) Explain why this cooling is not exponential.
(c) Calculate the initial rate of temperature fall.

Question 27

A patient lies in a hospital room at 16°C. The patient's skin is at 34°C, total skin area 1.6m², and emissivity 0.80.

(a) Find the net rate of radiative heat loss.
(b) A blanket with emissivity 0.05 is provided. Estimate the new net radiative heat loss from the blanket's outer surface, assuming the blanket surface reaches 22°C.

Question 28

A blackbody copper sphere (radius 30mm, density of copper $\rho_{\text{Cu}} = 8900\text{kgm}^{-3}$, specific heat capacity $c = 400\text{Jkg}^{-1}\text{K}^{-1}$) is cooled to −73°C and placed in an enclosure at 27°C. The initial rate of temperature rise is measured as 0.012Ks^{-1}. Use this data to calculate the Stefan-Boltzmann constant σ.

Question 29

A tungsten wire of length 100cm and perimeter 0.20cm operates at 2000K with emissivity 1.0. Find the rate of radiation from the wire, and the wavelength of maximum emission.

Question 30

A furnace radiates with maximum intensity at a wavelength of 1200Å. A star emits its maximum intensity at 4800Å.

(a) Find the temperature of each.
(b) Calculate the ratio of the power radiated per unit area by the furnace to that by the star.
(c) At what wavelength does the human body (37°C) radiate maximum intensity?

Question 31

An ice layer 10cm thick has formed on a pond. The air above is at −10°C and the water below is at 0°C. The thermal conductivity of ice is $k = 2.1\text{Wm}^{-1}\text{K}^{-1}$, the density of ice is $\rho_{\text{ice}} = 900\text{kgm}^{-3}$, and the latent heat of fusion is $L_f = 3.34 \times 10^5\ \text{Jkg}^{-1}$.

(a) Show that the time for the ice to grow from thickness x_1 to x_2 is given by $t = \frac{\rho_{\text{ice}} L_f}{2k\Delta T}(x_2^{\,2} - x_1^{\,2})$.

(b) Hence compare the time for the ice to grow from 10cm to 11cm with the time for it to grow from 1cm to 2cm.

Question 32

A hot-water pipe of inner radius 2.0cm and outer radius 3.0cm (copper, $k = 380\text{Wm}^{-1}\text{K}^{-1}$) and length 5.0m is lagged with a 4.0cm-thick layer of asbestos ($k = 0.08\text{Wm}^{-1}\text{K}^{-1}$). Water inside is at 80°C and the outer lagging surface is at 22°C. Find the rate of heat loss through the lagged pipe and the temperature at the copper-asbestos boundary.

Question 33

Earth's albedo is approximately $\alpha = 0.30$. The solar constant is $\varphi_s = 1380\text{Wm}^{-2}$.

(a) Modify the radiative equilibrium derivation to include the albedo, and show that the equilibrium temperature becomes lower than the bare-blackbody value.

(b) The polar ice caps have higher albedo than the oceans they cover. If global warming melts the ice caps, explain what happens to the overall albedo and what this implies for the equilibrium temperature.

(c) Is the effect described in part (b) a stabilising (negative) or a destabilising (positive) feedback? Explain.

Chapter 4

KINETIC THEORY OF GASES

INTRODUCTION

In the two weeks since Kipanga burned his hand on a stainless steel spoon in the school kitchen, he had turned one small injury into a long career. The red stripe on his palm had faded within three days, but the story had grown more impressive every time he told it. The water had been hotter. The spoon had been heavier. A rumour had started, somewhere in Form Four, that the heat had briefly stopped his wristwatch, and Kipanga, who did not own a wristwatch, had not confirmed the rumour, but he had not denied it either. He had also begun performing, at every morning break, a slow-motion reenactment of the grab-and-yelp, complete with sound effects that improved week by week, in front of a small audience of Form Two boys who found him genuinely heroic. He had been trying, without much success, to make "Mkono wa Moto" stick as his new nickname. Kipute, who had been there during the actual quarter-second event, had listened in silence for two weeks, and her silence was running out.

She found him in the courtyard that Friday morning, in the middle of the reenactment, frozen in an exaggerated pain-face in front of three rapt Form Two boys. She waited for the sound effect to finish. She waited for the applause. She waited for Kipanga to bow, which he did. Then she spoke with the calm of a person who had been preparing her words for fourteen days.

"Kipanga. A challenge. Make this bicycle pump noticeably hot in less than two minutes. No fire. No sunlight. Nothing rubbing against anything but air."

"I accept."

"I have not told you the prize."

"I accept it anyway. Mkono wa Moto does not need details."

"If you succeed, I will call you ***Professor Kipanga*** *for the whole of tomorrow. Every sentence, every corridor, every canteen. No rolling of the eyes. Full seriousness. For one day. But unless you can also explain, in physics, why the pump got hot, the title is only temporary, and I can take it back at any moment."*

"That is not how bets work."

"This is not a bet, Kipanga. This is a chapter."

A pump was borrowed from the gardener. A flat-tyred bicycle appeared from behind the water tank. The three Form Two boys, now promoted to witnesses, arranged themselves in a respectful semicircle. Kipanga pumped. The first twenty strokes were a performance. The next forty were the first real physical work of his morning. By the sixtieth, Kipanga had forgotten the crowd, the bet, and possibly his own name. He looked less like a man demonstrating physics and more like a man being slowly defeated by a piece of workshop equipment. He stopped. He rested his palm flat on the pump barrel. He snatched it away.

"Aa! It is hot."

One of the Form Two boys announced the time with the seriousness of a chief judge at the Olympics. One minute, fifty-one seconds. Kipanga lifted both arms above his head in triumph and, forgetting himself completely, let out a loud shout of "MKONO WA MOTO!" which the three Form Two boys, finally, repeated back in unison. It was the first time the nickname had actually worked in fourteen days. Kipanga was having an excellent morning.

Kipute walked forward four slow, careful steps, stopped in front of him, and bowed. The bow was small, slow, and completely serious.

"Congratulations, Professor Kipanga. Now, Professor Kipanga, would Professor Kipanga be kind enough to explain to the rest of us, in Professor Kipanga's own words, exactly how Professor Kipanga made the pump hot? Or is Professor Kipanga's knowledge on this matter a mystery even to Professor Kipanga himself?"

The three Form Two boys looked at each other. The word "Professor" had just appeared seven times in a single paragraph, each time with complete seriousness, as if at a state ceremony. One of them sat down on the concrete. Kipanga's raised arms came slowly down to his sides.

"The pump is hot because I pumped it."

"Professor Kipanga pumped it," corrected Kipute. *"Even Form One could tell us that. But no fuel was added. No match was struck. The sun was not on it. And Professor Kipanga's own rules forbid friction, because the inside of the barrel is hot too, and nothing inside rubs against anything except air. So where, Professor Kipanga, did the heat come from?"*

Kipanga opened his mouth. He closed it. He opened it a second time, slower, and closed it again with the weary finality of a man who has looked into his own mind and found that, on this particular question, his own mind is not at home.

At that moment Mr. Akilikubwa walked out of the staff room with his morning tea. He had been watching from the window for the last three minutes and had heard the word Professor six times already, which he thought, as a teaching tool, was almost perfect.

"Good morning, Professor Kipanga. A good piece of experimental physics. But a title given to you by Kipute is not a gift. It is a position, and a position comes with duties. The first and only duty of a physics Professor is to explain. So where, Professor, did the heat come from?"

"I do not know, sir."

Mr. Akilikubwa nodded, picked up the pump by the handle, and began to walk toward the physics laboratory. *"Bring the bicycle. We are going to make the same thing happen again, slowly enough to understand. Professor Kipanga, you may walk beside me, as a Professor should. Mkono wa Moto, I am sorry to say, will have to wait."*

Chapter 2 introduced temperature, and taught us to trust a thermometer over our fingers. Chapter 3 asked how heat moves, and gave us conduction, convection, and radiation. Between them, those two chapters have told us what heat does and where it goes. They have not once told us what temperature actually is. This is the chapter where physics finally answers that question.

Welcome to Chapter 4. Your own bicycle pump has been lying to you all your life about what is actually going on inside it. It is time to meet the molecules that tell the truth.

THE MOLECULAR PICTURE

Every Tanzanian student arriving in Form Five already knows what a gas is made of. From O-Level you have met the three states of matter, the bouncing of molecules in Brownian motion, the spreading of a smell across a room by diffusion, and the empirical gas laws that link pressure, volume, and temperature. What you have not yet been asked to do is to connect those molecular pictures to the numerical laws. Boyle's law and Charles' law were presented to you as experimental facts: pressure falls as volume rises, volume rises as temperature rises. Nobody told you why. Nobody asked what the molecules themselves were doing to produce those relationships. That question is the whole of Chapter 4.

The scale of what we are about to describe is worth pausing over. A single mole of gas contains roughly 6.0×10^{23} molecules, the Avogadro number, which means one breath of air carries more molecules than there are grains of sand on every beach on Earth. Those molecules are not close-packed. At atmospheric pressure the average spacing between neighbours is about ten times the diameter of a molecule, so the gas is mostly empty space with tiny fast-moving particles ricocheting through it. Each molecule travels in a straight line until it strikes another molecule or a wall, and each such collision lasts an extraordinarily short time compared with the time between collisions. The bicycle pump Kipanga was hammering was not a container of some continuous smooth substance. It was a container of about 10^{22} air molecules, each one moving at something like the speed of a rifle bullet, colliding with the barrel walls billions of times per second, and pushing back on his piston the only way they knew how.

The task of this chapter is to turn that picture into arithmetic. From six simple assumptions about how those molecules behave, we will derive the pressure a gas exerts on the walls of its container. From that pressure formula combined with the ideal gas equation, we will discover something genuinely surprising: the average kinetic energy of a gas molecule depends on one thing and one thing only, the Kelvin temperature. That single result will finally tell us what temperature is. It will also explain why Kipanga's pump got hot, why a pressure cooker works, why hydrogen leaks out of Earth's atmosphere while oxygen stays, and why a bicycle tyre left in the Dar es Salaam sun reads a higher pressure in the afternoon than it did in the morning. Six assumptions. One derivation. A great deal of physics. The next section starts with the assumptions.

ASSUMPTIONS OF THE KINETIC THEORY OF GASES

We had our molecular picture in the previous section. Now physics demands arithmetic. To derive a formula for the pressure a gas exerts on its container, we cannot work with an arbitrary cloud of complicated molecules doing whatever they please. We need a simplified model, one that keeps the essential physics (many small particles, random motion, collisions with walls) and throws away every complication the real gas would otherwise introduce (molecules of different shapes, molecules that stick to each other, collisions that lose energy, molecules with measurable size). The simplified model is called the ideal gas, and it is defined by six assumptions.

The six assumptions

Assumption 1: *A gas consists of a very large number of identical molecules moving in random directions with a range of different speeds.*

Assumption 2: *The volume occupied by the molecules themselves is negligible compared with the volume of the container.*

Assumption 3: *Collisions between the molecules and the walls of the container, and between molecules themselves, are perfectly elastic.*

Assumption 4: *Intermolecular forces are negligible except during the brief instant of collision.*

Assumption 5: *Between collisions, each molecule moves in a straight line at constant velocity, in accordance with Newton's first law.*

Assumption 6: *The time spent in a collision is negligible compared with the time between collisions.*

Stating the six assumptions correctly gives the basic framework. But understanding each one, knowing why it is there, and knowing when it breaks down, provides a much deeper grasp of the topic. We now take each of the six in turn and ask two questions of it: why is this simplification made, and under what conditions does it fail? A later section will return to these failures and use them to explain the difference between an ideal gas and a real gas. What follows here is the foundation of that comparison.

1. A very large number of identical molecules in random motion

Why this assumption is made. The arithmetic that produces the pressure formula relies on averages. An average is only meaningful when the number of values being averaged is **extremely large**, otherwise the result would fluctuate from one instant to the next and no single pressure could be assigned to the gas. One mole contains roughly 6.0×10^{23} molecules, which makes the statistical averaging exact for all practical purposes. The requirement that the molecules be **identical** means every one has the same mass *m* and behaves by the same rules, so the mass can be pulled out of the averaging as a single constant. **Randomness** of direction is what gives the gas its isotropy: no direction is preferred, which allows the derivation in the next section to split the total kinetic energy equally among the three Cartesian directions, a step we will depend on heavily.

Where it breaks down. This assumption fails only in systems where the number of particles is small enough for statistical fluctuations to matter. Nanoscale devices, or the thinnest upper regions of the atmosphere where only a handful of molecules cross a region per second, can show such fluctuations. In any gas you will meet at advanced-level, this assumption is essentially exact.

2. Molecular volume negligible compared with container volume

Why this assumption is made. The derivation that follows treats each molecule as a mathematical point, so that its position is a single location and it can be anywhere in the box with equal probability. This gives the gas access to the full volume *V* of the container. If the molecules had appreciable size, the available volume would be less than V, because each molecule would exclude a small region around itself from being occupied by another molecule. That complication would ruin the clean derivation and force us to carry an awkward correction through every equation.

Where it breaks down. At very high pressure, molecules are forced close together, and the fraction of the total volume occupied by the molecules themselves is no longer negligible. Inside a commercial oxygen cylinder at 200 atmospheres, molecules are packed tightly enough that the available volume is noticeably less than V, and the measured pressure is higher than PV=nRT would predict. The van der Waals correction, which you may meet in chemistry, introduces a replacement volume $V - nb$ to account for this exclusion, where b is a constant fitted to each real gas.

3. Perfectly elastic collisions

Why this assumption is made. Elastic collisions conserve total kinetic energy. If collisions were inelastic, if molecules lost a little energy each time they struck the wall or each other, the gas would slowly cool to zero temperature without any heat being removed, which is plainly not what happens. A gas in thermal equilibrium, left alone in a sealed insulated container, keeps the same temperature forever. The only way the molecular model can reproduce that stubborn fact is to require every collision to conserve kinetic energy. Without this assumption, the pressure formula we will soon derive would decay with time, and so would the temperature.

Where it breaks down. No gas violates this in the conditions an A-level student will meet. A very precise treatment of real gases accounts for inelastic molecule-on-molecule collisions where some kinetic energy is briefly transferred into rotational or vibrational excitation of the colliding molecules. In polyatomic gases this energy exchange becomes important when we discuss degrees of freedom later in this chapter. For the pressure derivation itself, elastic collisions at the walls is all we need.

4. Negligible intermolecular forces except during collision

Why this assumption is made. If two molecules attracted each other from a distance, their motion between collisions would not be a straight line, they would curve toward each other. The pressure derivation would then need to track those curving paths, which is not possible in closed form. By assuming no forces act between molecules except during the instant of direct contact, we guarantee that every molecule travels in a straight line between collisions, which is what makes assumption 5 consistent and what keeps the derivation tractable.

Where it breaks down. At low temperatures, molecules move slowly enough that the small attractive forces between them, the same van der Waals forces responsible for the volume correction above, become non-negligible. The molecules tend to linger near each other in brief clusters. This is exactly what allows real gases to be liquefied at low enough temperatures: an ideal gas, by assumption, cannot be liquefied at any pressure because it has no attractive forces to pull the molecules together.

A real gas can. This is why the last section of this chapter will state that ideal behaviour requires both low density and high temperature.

5. Straight-line motion between collisions

Why this assumption is made. This is a direct consequence of assumption 4 plus Newton's first law. A molecule with no force acting on it moves in a straight line at constant velocity. The derivation in the next section tracks a single molecule between two consecutive collisions with the same wall, and that tracking requires the molecule to travel a known distance (a round trip of 2L in a cubic box of side L) in a known time (2L/u where u is its x-component of velocity). A molecule that was curving between collisions would travel further than 2L between hits, and the time between hits would not be given by our simple formula.

Where it breaks down. Same physics as assumption 4, because it is the same physics. Attractive forces curve the paths slightly at close range; at low temperature and high density those curves matter.

6. Collision time negligible compared with time between collisions

Why this assumption is made. During the brief instant of a wall-collision, the molecule is in contact with the wall and is experiencing a large force, it is not behaving as a free particle. During the much longer time between collisions, it is a free particle obeying Newton's first law. The derivation treats collisions as instantaneous events that reverse the x-component of velocity and leave everything else unchanged. This simplification is valid provided the molecule spends an overwhelmingly larger fraction of its time travelling freely than it does in contact with the wall. In air at room temperature, a typical molecule travels for something like 10^{-9}s between collisions and each collision itself lasts only about 10^{-13}s, so the ratio is ten thousand to one in favour of free motion. The assumption is excellent for all ordinary conditions.

Where it breaks down. At pressures so high that molecules collide almost continuously (liquids, dense gases, the interior of a giant planet), this assumption fails. The molecules are effectively always in contact, and the free-flight picture no longer applies. This is one of several reasons why liquids are not described by PV=nRT.

These six assumptions together define what physicists call an **ideal gas**. An ideal gas is not a real substance. It is a deliberately idealised model in which molecules are vanishingly small, never interact except on contact, and never lose energy. Real gases behave like ideal gases only when three conditions are met together: **low pressure** (so molecular volume is negligible compared to container volume), **high temperature** (so molecules move quickly enough that attractive forces between them do not matter), and **low density** (so molecules collide rarely enough that the time-between-collisions assumption is comfortable). The last section of this chapter will return to these conditions and examine what happens when they are not met. For now, we take the six assumptions as given and put them to work: in the next section we will use them to derive the pressure of an ideal gas from first principles.

Before we leave this section, a small test. The Professor who won his title in the introduction is about to lose it to a planet. Listen carefully.

REAL Example 1

Kipute approaches Mr. Akilikubwa in the staff room during break, with a physics textbook open at a picture of the planet Jupiter.

Kipute: *"Sir, this chapter says Jupiter has a gaseous atmosphere going all the way down to its core, and that the pressure near the centre is millions of atmospheres. Is it still a gas at those pressures?"*

Mr. Akilikubwa: *"Technically, yes. The material is gaseous in the sense that its molecules are not locked into a crystal. But it behaves nothing like the air in this room. Tell me which of our six assumptions would fail first if we tried to apply PV=nRT to the centre of Jupiter."*

Kipute (thinking): *"At millions of atmospheres, molecules are crushed very close together. The spacing between them must be tiny. So the volume occupied by the molecules themselves is no longer small compared with the volume they are confined to. That is assumption 2."*

Mr. Akilikubwa: *"Correct. What else?"*

Kipute: *"When molecules are that close, they must be in contact nearly all the time. The attractive forces between them would be significant. That is assumption 4."*

Mr. Akilikubwa: *"Also correct. And the consequence?"*

Kipute: *"Assumption 5 then fails automatically, because if forces are acting between collisions, the molecules do not move in straight lines. And if molecules are essentially always in contact, assumption 6 also fails, because there is no free-flight time between collisions."*

Mr. Akilikubwa: *"So how many of our six fail at high density?"*

Kipute: *"Four. Two, four, five, and six."*

Mr. Akilikubwa: *"Exactly four. Only assumptions 1 and 3 survive. The number of molecules is still large, and collisions are still elastic to the same degree they always are. This is why gas giants are not described by PV=nRT and why the equation we are about to derive is useful only in the limit of low density and high temperature, which, thankfully for the NECTA examiners and for this chapter, describes every gas you will ever be asked to analyse at Form Five level."*

Making Sense of the Answer: *The point of this dialogue is not Jupiter. The point is that the six assumptions are not independent of one another. When one fails at high density, several others fail at the same time, because they describe different aspects of the same underlying physics, namely that molecules must be widely spaced and weakly interacting for the ideal gas model to apply. Remember this pattern. It will return in the final section of this chapter when we ask under what conditions a real gas behaves ideally, and we will discover that the answer is essentially an inversion of the Jupiter case: low pressure, high temperature, low density.*

Think Like a Physicist: *When an exam question asks which of the six assumptions fails under some extreme condition, do not evaluate all six one by one. Find the single physical quantity that the condition is pushing, high density, low temperature, small particle count, and then ask which assumptions are built on that quantity being extreme. The assumptions cluster. High density always breaks 2, 4, 5, and 6 together, because all four of them silently assume molecules are widely spaced and independent between contacts. Low temperature always breaks 4 and 5 together, because both depend on attractive forces being weak. Learn the clusters, not the individual failures, and you will answer these questions in half the time.*

PRESSURE OF AN IDEAL GAS

The six assumptions in the previous section give us everything we need to start calculating. We are about to derive the most important result in the kinetic theory of gases, a formula that tells us, from first principles and starting from nothing more than Newton's laws, what pressure a gas of many molecules will exert on the walls of its container. The derivation will follow the career of a single molecule through one round trip inside a cubical box, use Newton's second law to turn the molecule's wall-collisions into a force, extend the result to all the molecules in the container, and close with a symmetry argument that connects the motion along one axis to the motion in all three. Every stage will rest on one or more of the six assumptions. By the end, you will understand not just the final formula but what each assumption is doing along the way to keep it honest.

Consider a cubical container of side **L** containing a single gas molecule of mass **m**. The molecule is moving in some arbitrary direction with some speed, but for the moment we will focus only on the component of its motion along the x-axis, which we call **u**. The figure below shows the geometry. The molecule bounces back and forth between the two walls perpendicular to the x-axis, the wall at x = 0 on the left and the wall at x = L on the right, and we are going to calculate how much force it exerts on one of those walls, say the right one.

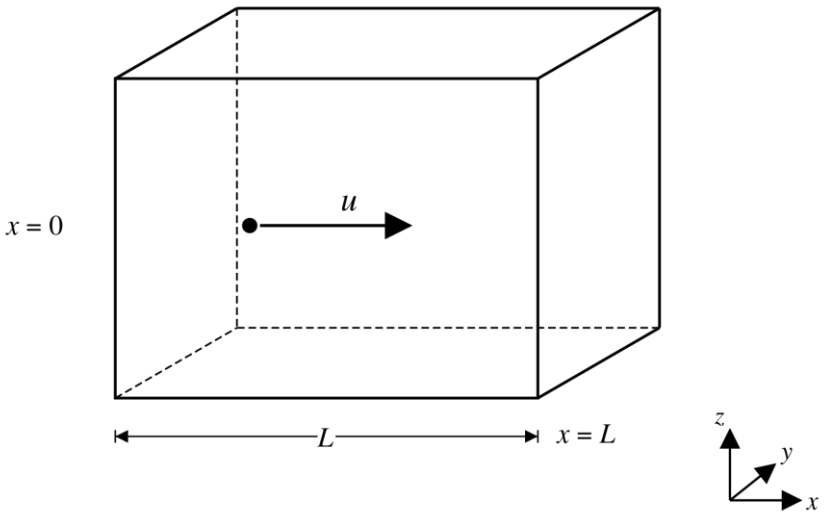


Figure: *A single gas molecule of mass m inside a cubical container of side L, moving with x-component of velocity u. The molecule bounces elastically between the walls at x = 0 and x = L.*

Determining the momentum change in a single collision

Consider one bounce at the wall x = L. Just before the bounce, the molecule's x-component of momentum is +mu. Just after the bounce, because collisions with the walls are perfectly elastic (assumption 3), the x-component of velocity has reversed

to −u, and so the momentum is now −mu. The y- and z-components of the molecule's momentum are unchanged by this collision, because the wall acts only along the x-direction. The change in the x-component of momentum is therefore:

$$\Delta(\text{mu}) = (-\text{mu}) - (+\text{mu}) = -2\text{mu}$$

The magnitude of this change is 2mu. By Newton's third law, the wall has received an equal and opposite impulse of +2mu on every such bounce. Every time the molecule strikes the wall, it delivers a little kick of size 2mu to it. The question we must answer next is: *how often does this kick arrive?*

Determining the time between collisions

Between successive bounces the molecule travels in a straight line at constant velocity (assumptions 4 and 5). Starting from the wall at x = L it moves off to the opposite wall at x = 0, a distance L, reverses there (another elastic bounce), and returns to the wall at x = L, another distance L. The total round-trip distance between one hit on the right-hand wall and the next hit on the same wall is therefore 2L. At constant x-speed u, the time this round trip takes is:

$$\text{t} = \frac{2\text{L}}{\text{u}}$$

This is the time between successive kicks delivered by our single molecule to the wall at x = L.

Determining the average force on the wall

We now have everything we need for Newton's second law. In its most general form, Newton's second law states that force is the rate of change of momentum. Our molecule delivers momentum of size 2mu to the wall every time 2L/u, so the average force the molecule exerts on the wall is:

$$\text{F} = \frac{2\text{mu}}{\frac{2\text{L}}{\text{u}}} = \frac{\text{mu}^2}{\text{L}}$$

This is the average force that one molecule exerts on one wall, expressed purely in terms of the molecule's mass, its x-speed, and the size of the container.

A word of care about what this "average" really means. The molecule is not pressing on the right-hand wall continuously. It strikes the wall once per round trip, delivers its impulse in an instant, then spends the next 2L/u seconds flying away toward the opposite wall, returning, and preparing to strike again. Between consecutive hits, its contribution to the force on the right-hand wall is exactly zero. The formula $F = mu^2/L$ gives the time-averaged force over one complete round trip, smoothing this on-off delivery into a single equivalent steady force. It is this round-trip-average, and not an instantaneous force, that we will add up when we sum over all molecules in the next step. Every molecule contributes the same kind of round-trip-averaged force to every wall it can reach, and the bookkeeping works out cleanly precisely because each molecule spends equal time heading toward, and equal time heading away from, any given wall.

Determining the total force on the wall from all N molecules

A real gas does not contain a single molecule. It contains an enormous number, N of them (assumption 1). Each molecule has its own x-component of velocity: call them u_1, u_2, and so on up to u_N. Because the molecules do not interact except during instantaneous collisions (assumption 4), and because those collision times are negligible (assumption 6), each molecule contributes its own round-trip-averaged force to the right-hand wall independently. The total average force on the wall is simply the sum of all the individual contributions:

$$\text{F}_{\text{total}} = \frac{\text{m}(\text{u}_1^2 + \text{u}_2^2 + \ldots + \text{u}_\text{N}^2)}{\text{L}}$$

This sum of squares cries out for averaging. We define the **mean square** of the x-components as *the average of the individual squared velocities*:

$$\bar{\text{u}}^2 = \frac{\text{u}_1^2 + \text{u}_2^2 + \ldots + \text{u}_\text{N}^2}{\text{N}}$$

Multiplying both sides by N gives the original sum of squares as $N\bar{u}^2$, and substituting this into the force expression above gives the total force on the wall in its final compact form:

$$\text{F}_{\text{total}} = \frac{\text{Nm}\bar{\text{u}}^2}{\text{L}}$$

Determining the pressure on the wall

The right-hand wall is a square face with area L^2. Pressure is the force per unit area of wall, averaged over time and over the whole surface. Dividing the total force above by the wall area, and recognising that $L^3 = V$ is the volume of the container:

$$\text{P} = \frac{\text{F}_{\text{total}}}{\text{L}^2} = \frac{\text{Nm}\bar{\text{u}}^2}{\text{L}^3} = \frac{\text{Nm}\bar{\text{u}}^2}{\text{V}}$$

Notice that the division by wall area only becomes physically legitimate now, at the end of the sum. A single molecule does not spread its force over L^2 of wall. It strikes a single point and moves on. Pressure is a collective quantity: it has physical meaning only once we have added the contributions from many molecules whose collision points are scattered randomly across the whole wall. Assumption 1, random motion, is what guarantees that the scattering is uniform and that the total force really is distributed over the full L^2.

At this stage the formula is exact. But it is not yet useful, because it has been written in terms of the mean square of the x-components of velocity alone, and the x-direction is not special. Anyone setting up the container with the axes labelled differently would have produced the same formula with the mean of $\bar{u}^2$ replaced by $\bar{v}^2$ or $\bar{w}^2$. The molecules move in three dimensions, and what we really want is a formula in terms of the full three-dimensional motion. The next step makes that connection.

Invoking isotropy

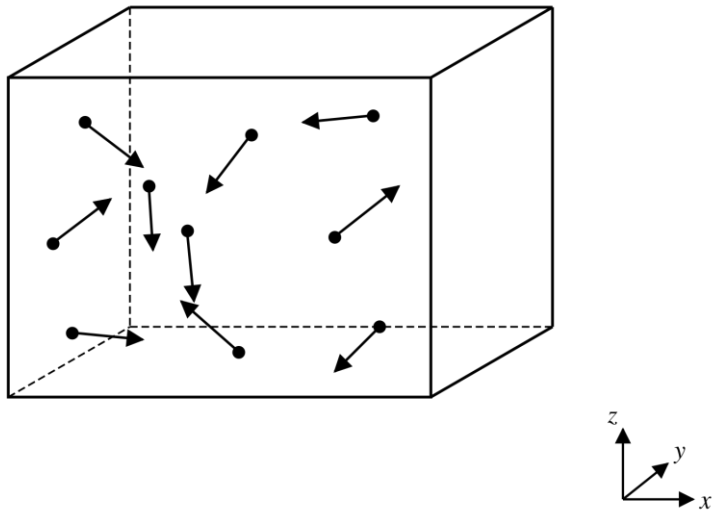


Figure: *The same cubical container, with ten representative molecules drawn out of the enormous number present. Their velocity arrows point in every direction. No direction is preferred. This randomness is the engine of the isotropy argument.*

For any individual molecule, let u, v, w be its three velocity components along the x-, y-, and z-axes, and let c be the magnitude of its resultant velocity. By the ordinary Pythagorean rule in three dimensions:

$$c^2 = u^2 + v^2 + w^2$$

This equation is true for every single molecule in the container. Taking the average of both sides over all N molecules (*the average of a sum is the sum of the averages*):

$$\bar{c}^2 = \bar{u}^2 + \bar{v}^2 + \bar{w}^2$$

Now we invoke the full force of assumption 1. The molecules are moving in completely random directions. No direction of space is preferred, neither the x-axis nor the y-axis nor the z-axis has any special claim on the motion. The average of any quantity built out of the x-component of velocity must therefore be identical to the average of the corresponding y- and z-quantities:

$$\bar{u}^2 = \bar{v}^2 = \bar{w}^2$$

Picture it. *Inside the cube, the molecules are flying simultaneously in every imaginable direction: towards the ceiling, towards the floor, towards the left wall, towards the right wall, towards every diagonal in between. No direction gets a special invitation. On average, exactly one-third of each molecule's kinetic energy is associated with its motion along the x-axis, exactly one-third with its motion along the y-axis, and exactly one-third with its motion along the z-axis. This is not an approximation, not an estimate, not a rough rule of thumb. It is an exact consequence of the randomness we assumed at the very beginning. Randomness gives us isotropy, and isotropy gives us the factor of one-third that is about to appear in the final formula.*

With the three mean squares equal, the sum on the previous line reduces immediately to $\bar{c}^2 = 3\bar{u}^2$, or rearranged:

$$\bar{u}^2 = \frac{\bar{c}^2}{3}$$

The central formula

Substituting this into the pressure formula $P = Nm\bar{u}^2/V$ gives the central result of this section:

$$\mathbf{P = \frac{Nm\bar{c}^2}{3V}}$$

or, multiplying both sides by V:

$$\mathbf{PV = \frac{1}{3}Nm\bar{c}^2}$$

This is the pressure of an ideal gas expressed in terms of its molecular properties. The equation is called the **kinetic equation of gases**, and everything else we will do with the molecular picture in this chapter flows from it.

The same result can be rewritten in a form that is usually more convenient in problems. The total mass of gas in the container is Nm (the number of molecules times the mass of each one), and the volume of the container is V, so the density of the gas is:

$$\rho = \frac{Nm}{V} \text{ or } Nm = \rho V$$

Substituting $Nm = \rho V$ into the pressure formula:

$$\mathbf{P = \frac{\rho \bar{c}^2}{3}}$$

It is customary to write the mean square speed $\bar{c}^2$ as v_{rms}^2, where v_{rms} is called the **root mean square speed** of the molecules. In that notation:

$$\mathbf{P = \frac{1}{3}\rho v_{rms}^2}$$

which rearranges to give v_{rms} directly from the pressure and density of the gas:

$$\mathbf{v_{rms} = \sqrt{\frac{3P}{\rho}}}$$

These three formulas are all the same physics rewritten in different clothing. Which one you reach for in a given problem is a matter of which quantities you have been given. If you know the number of molecules and the mass of each, use the first. If you know the density and the pressure, use the third to find the rms speed in one line. Nobody will ask you to reproduce the whole derivation on paper in an examination, but every formula in this summary is examinable, and most problems in this chapter will use at least one of them. Six examples follow, each exercising a different muscle.

BINDER Example 2

A sealed cubical container of volume 5.0L holds 3.0×10^{24} molecules of gas, each of mass 3.3×10^{-26}kg. The root mean square speed of the molecules is 400m/s. Calculate the pressure of the gas inside the container.

Solution

Three of the molecular properties are given to us directly: N, m, and v_{rms}. The volume is given in litres and must be converted to cubic metres before substitution. The formula of choice is therefore:

$$P = \frac{Nmv_{rms}^2}{3V}$$

Converting the volume:

$$V = 5.0\,L = 5.0\times10^{-3}\,m^3$$

Substituting the numerical values:

$$P = \frac{(3.0\times10^{24})(3.3\times10^{-26}\,kg)(400\,m/s)^2}{3\times5.0\times10^{-3}\,m^3} = 1.056\times10^6\,Pa \approx 1.06\,MPa$$

The pressure inside the container is 1.056×10^6 Pa (or about 1.06MPa, roughly ten times atmospheric pressure).

Making Sense of the Answer: *Ten atmospheres is the kind of pressure you would find in a fully inflated bicycle tyre, or in a small compressed-gas canister. It is comfortably within the range of pressures you would expect from a sealed container holding something like five moles of gas (3.0×10^{24} molecules is about five moles) in only five litres. The molecules are more tightly packed than they would be at atmospheric pressure, so they hit the walls more often per second, so the pressure is higher. Everything adds up.*

Think Like a Physicist: *The two most common slips in a calculation of this kind are forgetting that $1L = 10^{-3}\,m^3$ (a factor of a thousand), and forgetting that v_{rms} appears in the formula squared, not to the first power. Both errors are silent until the final answer, where they produce results that are wrong by factors of thousands or by large factors of 400. Write the unit after every number as you substitute, and watch them multiply and cancel. If the units at the end come out as Pa, you have almost certainly done the calculation right.*

BINDER Example 3

At standard temperature and pressure the density of nitrogen gas is 1.25kg/m³. Calculate the root mean square speed of the nitrogen molecules.

Solution

The third summary formula was built precisely for this case: the pressure and density are given, and the rms speed is the unknown:

$$v_{rms} = \sqrt{\frac{3P}{\rho}}$$

At standard temperature and pressure: T=273K, P=101325Pa
Substituting the numerical values:

$$v_{rms} = \sqrt{\frac{3 \times 101325\text{Pa}}{1.25\text{kg/m}^3}} = \sqrt{243180\text{m}^2/\text{s}^2} = 493.1\text{m/s}$$

The root mean square speed of the nitrogen molecules at STP is about 493m/s.

Making Sense of the Answer: *493m/s is roughly 1800km/h, about 1.4 times the speed of sound in air (which is close to 340m/s at this temperature). That ordering is universal: the rms speed of a gas always exceeds the speed of sound in that same gas, because a sound wave is a coordinated, cooperative vibration of many molecules moving together, while v_{rms} measures the uncoordinated individual speeds in all directions. A sound wave can only travel as fast as the molecules can physically carry the disturbance, and they cannot carry it faster than they themselves are moving. The factor of roughly 1.4 between v_{rms} and the speed of sound is not accidental, it turns out to equal $\sqrt{(3/\gamma)}$ where γ is a ratio we will meet in a later chapter. For now, note only that the two speeds are of the same order of magnitude, with v_{rms} always the larger of the two.*

Think Like a Physicist: *The unit reduction Pa/(kg/m³) → m²/s² is worth learning by heart. It is the reason the formula $v_{rms} = \sqrt{(3P/\rho)}$ returns a speed in m/s when the pressure is in pascals and the density is in kilograms per cubic metre, with no extra conversion factors. Every time you meet a formula whose units work out cleanly on the first try, take a moment to notice. SI was designed that way for exactly this convenience.*

BINDER Example 4

A balloon contains helium gas at a pressure of 1.2×10^5Pa, and the helium has a density of 0.18kg/m³.

(a) Calculate the root mean square speed of the helium atoms.

(b) The balloon is gently squeezed at constant temperature to half its original volume. Find the new pressure inside the balloon, and explain the reasoning in terms of what stays the same and what changes at the molecular level.

Solution

(a) Using $v_{rms} = \sqrt{(3P/\rho)}$:

$$v_{rms} = \sqrt{\frac{3 \times 1.2 \times 10^5\ \text{Pa}}{0.18\ \text{kg/m}^3}} = 1414\ \text{m/s}$$

The rms speed of helium atoms at this pressure and density is about 1414m/s (which is roughly 1.41km/s).

(b) At constant temperature, the average kinetic energy per molecule does not change (this is a result we will prove formally in a later section of this chapter, but for now take it as given). Since the mass of each helium atom is fixed, the rms speed must also stay the same: the molecules do not speed up or slow down when the balloon is squeezed slowly at constant temperature.

What does change is the density. The same mass of helium now occupies only half the volume, so the density has doubled. From the formula $P = (1/3)\rho v_{rms}^2$, with the v_{rms}^2 unchanged and the density doubled, the pressure must also double:

$$P_{new} = 2 \times 1.2 \times 10^5\ \text{Pa} = 2.4 \times 10^5\ \text{Pa}$$

The new pressure is 2.4×10^5Pa, exactly twice the original.

This is Boyle's law, which you met at O-level as an empirical fact, now derived from the molecular picture: *at constant temperature, halving the volume doubles the pressure because it doubles the density without changing the molecular speeds.*

Making Sense of the Answer: *The helium rms speed (about 1414m/s) came out much larger than the nitrogen rms speed in the previous example (493m/s at nearly the same pressure). This is not because the helium is somehow hotter or more energetic; the two gases in these two problems are at similar conditions. Helium moves faster because its atoms are lighter: a helium atom is about seven times lighter than a nitrogen molecule, so for the same average kinetic energy per molecule, the lighter helium atom must move faster by a factor of $\sqrt{7} \approx 2.6$, which is close to the ratio $1414/493 \approx 2.87$ that we actually observe. The small discrepancy comes from slightly different temperatures in the two problems. You will see this mass-dependent speed pattern over and over in this chapter: lighter molecules move faster at the same temperature. It is the reason hydrogen leaks out of Earth's atmosphere but oxygen does not.*

Think Like a Physicist: *Whenever a gas-law problem says "at constant temperature", do not reach for the calculator first. Ask: what stays the same (temperature → kinetic energy per molecule → v_{rms}), what changes (volume → density), and what does the pressure formula then tell you. Three-quarters of Boyle's-law problems are solved in the head before you pick up a pen. The remaining quarter is just checking that the arithmetic matches what the physical reasoning has already told you must happen.*

REAL Example 5

A deflating balloon with a small pinhole emits a steady hiss and pushes a steady jet of gas out of the hole for several seconds, yet the kinetic theory states that the pressure of the gas inside the balloon is produced by discrete molecular collisions with the inner wall, each one a sharp, brief, isolated impulse at a single point. Explain, in terms of the molecular picture, why the pressure driving the gas through the hole is steady rather than erratic, naming the two averaging processes responsible.

Solution

The pressure driving the gas through the hole is not the force of a single collision but the time-averaged and area-averaged force delivered to the balloon wall surrounding the hole. Two averaging processes operate together.

1) **Averaging in time:** many millions of molecular collisions arrive at every patch of wall during each millisecond, so the running sum of those collisions over any reasonable measurement interval is smooth.
2) **Averaging in area:** the wall is struck by molecules across its whole surface, and by the randomness (of assumption 1) the collision points are distributed uniformly, so the force per unit area is the same everywhere.

So the steady pressure driving the steady jet is the doubly averaged result of enormously many microscopic events, and the hiss is steady for the same reason: the listener's ear averages thousands of tiny pressure fluctuations per second into a single perceived continuous sound.

Making Sense of the Answer: *Pressure, temperature, density, and every other macroscopic thermal quantity in this book behaves like a smooth continuous variable even though the molecular reality beneath it is grainy and discrete. The smoothness is never the bare physics; it is the physics plus the averaging properties of whatever instrument, or organ, or thin wall of rubber, is sampling the molecular events.*

Think Like a Physicist: *Every time a macroscopic measurement appears smooth and steady, ask what microscopic process is being averaged to produce the smoothness and over what timescale and length scale the averaging is happening. A faster instrument, or one with a smaller sensitive face, would begin to see the fluctuations that a slower instrument hides. What counts as "smooth" is as much a property of the measurement apparatus as it is of the underlying physics.*

HOT Example 6

A bicycle pump barrel is sealed at the far end and open at the near end, where a piston slides freely. A pressure gauge is fitted to the sealed end. With the piston fully withdrawn, the sealed air column occupies the full length L of the barrel and the gauge reads 2.0×10^5Pa. **Kipanga** pushes the piston in **gently** and **slowly**, and holds it at a position where the sealed air column has shrunk to one third of its original length. Throughout the compression, the air inside the barrel remains at steady room temperature.

(a) Predict the new gauge reading.
(b) Show by direct analysis that the per-molecule force has tripled compared with its value when the piston was fully withdrawn.
(c) Kipanga, in the chapter introduction, pumped his bicycle tyre vigorously for nearly two minutes and observed that the pump barrel became noticeably hot. If Kipanga had compressed this same pump barrel quickly rather than gently and slowly, would the gauge reading at the one-third-length position be equal to the answer of part (a), larger than it, or smaller? Identify which assumption of the present problem would be violated, and explain qualitatively why.

Solution

(a) At constant temperature, the average kinetic energy per molecule does not change, so the mean square speed $\bar{c}^2$ is unchanged. The pump is sealed, so the number of molecules N is unchanged. What changes is the volume of the sealed air column: the cross-sectional area A of the barrel is fixed, so the volume drops from AL to $A\left(\frac{L}{3}\right) = \frac{AL}{3} = \frac{V}{3}$. From the kinetic equation $P = \frac{Nm\bar{c}^2}{3V}$, with N, m, and $\bar{c}^2$ all unchanged and V replaced by one-third its original value, the pressure must triple:

$$P_{new} = \frac{Nm\bar{c}^2}{3\times\frac{V}{3}} = 3\times\frac{Nm\bar{c}^2}{3V} = 3P$$

$$P_{new} = 3 \times 2.0\times10^5 \text{ Pa} = 6.0\times10^5 \text{ Pa}$$

The gauge reading at the one-third-length position is 6.0×10^5Pa, three times the original, which is the molecular derivation of the same result Boyle's law would give directly.

(b) A single molecule with x-component of velocity u exerts a round-trip-averaged force on the sealed end equal to:

$$F = \frac{mu^2}{L'}$$

where L' is the length of the sealed air column, the distance the molecule must travel between successive hits on the sealed end. Initially, with the piston fully withdrawn, $L' = L$, so the initial per-molecule force is:

$$F_1 = \frac{mu^2}{L}$$

After compression, $L' = L/3$, so the new per-molecule force is:

$$F_2 = \frac{mu^2}{L/3} = \frac{3mu^2}{L} = 3F_1$$

Hence the per-molecule force has tripled.

(c) If Kipanga compresses the barrel quickly rather than gently, the gauge reading at the one-third-length position will be **larger** than the 6.0×10^5Pa predicted by part (a).

The assumption being violated is that the temperature of the gas inside the barrel stays constant during the compression. In the idealised slow compression of parts (a) and (b), heat has time to flow out through the barrel walls into the surrounding air, keeping the gas at room temperature throughout. In a fast compression, there is no time for that heat flow. The mechanical work done by Kipanga's hand on the piston is pushed into the gas and cannot escape fast enough, so the temperature of the gas rises during the compression. A higher temperature means a larger mean square speed, and therefore, by the kinetic equation $P = Nm\bar{c}^2/(3V)$ with $\bar{c}^2$ now increased as well as V decreased, the pressure rises by more than the Boyle's-law factor of three. (This is exactly the physics that heated the pump barrel in the chapter introduction and briefly made Kipanga the proud owner of the title Mkono wa Moto. A slow, patient compression converts work into pressure alone. A fast, vigorous compression converts work into pressure and heat together.)

Making Sense of the Answer: *The two parts of this example have shown the same physics from two different angles. Part (a) used the finished formula $P = Nm\bar{c}^2/(3V)$ and simply noted that shrinking V by a factor of three triples P. Part (b) went all the way back to derivation of the kinetic equation of gases and showed, molecule by molecule, exactly why that factor of three appears. The reason is the round-trip distance: shortening it makes each molecule hit the sealed end more often per second, not harder per hit. It is instructive that nothing about the force of a single collision changed. The molecules are still arriving at the same 2mu per kick they always delivered. What changed was the rate of arrivals, because the molecules now have a shorter commute between hits.*

Think Like a Physicist: *Parts (a) and (b) did not compute the same quantity. Part (a) found a gauge reading; part (b) found a per-molecule force. What they share is a scaling factor: both the collective pressure and the single-molecule force triple when the column shortens to one third. This is the useful habit to carry forward. Whenever a macroscopic result surprises you, ask what it predicts at the single-molecule level and check the two agree on ratios, not on numbers. Ratios are the currency of consistency between the microscopic picture and the macroscopic one.*

HOT Example 7

Nitrogen gas at standard temperature and pressure (273K, 1.013×10^5Pa) has density 1.25kg/m³ and a root mean square speed of approximately 493m/s. Consider a flat region of container wall of area 1.00m² in contact with this gas.

(a) Calculate the total average force delivered to this one square metre of wall by the gas.
(b) The typical x-component of molecular velocity, averaged across only those molecules currently heading toward the wall, is approximately $v_{rms}/\sqrt{3}$. Use this to estimate the momentum kick delivered to the wall by one average molecule in a single collision, given that the molecular mass of N_2 is 4.65×10^{-26}kg.
(c) Estimate the number of molecular collisions striking the one-square-metre region of wall each second. Comment on the size of your answer and what it tells you about why the pressure appears steady.

Solution

(a) Pressure is force per unit area, so the total average force on a one-square-metre region of wall is simply:

$$F = P \times A = 1.013\times10^5\ \text{Pa} \times 1.00\ \text{m}^2 = 1.013\times10^5\text{N}$$

The one square metre of wall receives, on average, 1.013×10^5 newtons of force from the gas at every instant, which is roughly the weight of ten tonnes pressing on it. (This is actually the same force on both sides, which is why the wall does not collapse inward; atmospheric pressure presses on the outside equally hard.)

(b) Using the given v_{rms} = 493m/s:

$$u_{avg} = \frac{493\ \text{m/s}}{\sqrt{3}} = 284.6335\text{m/s}$$

The momentum change delivered to the wall in one collision is 2mu, so the average momentum kick per collision is:

$$\Delta p = 2mu_{avg} = 2 \times 4.65\times10^{-26}\ \text{kg} \times 284.6335\text{m/s} = 2.65\times10^{-23}\ \text{kgm/s}$$

Each molecule arriving at the wall delivers about 2.65×10^{-23} kgm/s of momentum in a single collision (an extraordinarily small number).

(c) The total force on the wall is the total rate of momentum delivery: force equals momentum per second. So, if n is the number of collisions per second on the one-square-metre region, and each collision delivers on average Δp of momentum, then:

$$F = n \times \Delta p$$

Rearranging and substituting:

$$n = \frac{F}{\Delta p} = \frac{1.013 \times 10^5 \text{ N}}{2.65 \times 10^{-23} \text{ kgm/s}} \approx 3.8 \times 10^{27} \text{ collisions per second}$$

Comment: the number of collisions per square metre per second is of order 10^{27}, which is extremely large. Each collision delivers a very small momentum, but their combined rate is so large that the total force per unit area is smooth and constant, which explains why the pressure of the gas appears steady even though it is produced by discrete molecular events.

Making Sense of the Answer: *The number 3.8×10^{27} per square metre per second is not a unit to be memorised; it is a picture to be kept in mind whenever the word "pressure" appears in a physics problem. What we casually call the pressure of a gas is, at the molecular level, a storm. It is four octillion separate collision events pounding every square metre of container wall every second, each one too small and too brief to be noticed individually, their combined effect smooth and steady only because there are so enormously many of them. The REAL Example 5 about the balloon, explained why this storm looks steady to a pressure gauge. This calculation has shown how intense the storm actually is. Together the two examples close out this section with a single message: pressure is not what it looks like. It is what emerges when you average.*

Think Like a Physicist: *Reversing a formula is one of the most useful and least taught skills in physics. The pressure formula P = F/A is usually reached for when pressure is unknown; this example reached for it when the collision rate was unknown. Every formula in physics can be rearranged to answer a different question than the one it was originally introduced to answer. When you meet a new formula, make a habit of asking: if I knew every quantity in this equation except one, which quantities could be the unknown, and what physical question would each unknown correspond to? You will find that most formulas secretly answer half a dozen different questions depending on which variable you are solving for, and the richer your repertoire of reversals, the more problems you will find you already know how to solve.*

We have now derived the central formula (called kinetic equation of gases) of this chapter and tested it against seven different kinds of problem. The pressure of an ideal gas is set by its density and the mean square speed of its molecules. What we have not yet done, and what part (c) of HOT Example 6 had to borrow in advance, is show that the average kinetic energy per molecule depends only on the Kelvin temperature. For that we need to bring in the ideal gas equation and compare it with what we have just derived. That is the work of the next two sections.

THE IDEAL GAS EQUATION

The kinetic equation we have just derived is a beautiful piece of physics. It tells us how the pressure of a gas arises from the rain of molecular collisions on the container walls, and it gives us a formula for that pressure in terms of the number of molecules, their mass, and their mean square speed. But it has one limitation that makes it awkward for practical problems: it does not contain the temperature of the gas anywhere.

What we need is a second equation. An equation that speaks the language of pressure, volume, temperature, and the amount of gas you can measure out of a bottle. This second equation must be consistent with the kinetic equation; the two of them, placed side by side, must describe the same gas and produce answers that agree with each other. And when we put them side by side, something quietly spectacular will happen, but we will keep that surprise for the end of this section.

The second equation is called the **ideal gas equation**, and it comes not from the molecular picture but from three empirical laws that you met at O-Level. Each law describes how one pair of gas variables behaves when the others are held constant. We remind ourselves of all three, combine them into one, and give the resulting constant of proportionality a name. That is the whole plan for this section.

Three empirical gas laws

Boyle's law states that for a fixed mass of gas held at constant temperature, the volume is inversely proportional to the pressure:

$$V \propto \frac{1}{P} \qquad \text{(constant n and T)}$$

Charles' law states that for a fixed mass of gas held at constant pressure, the volume is directly proportional to the absolute temperature:

$$V \propto T \qquad \text{(constant n and P)}$$

Avogadro's law states that for a gas held at constant pressure and constant temperature, the volume is directly proportional to the amount of gas:

$$V \propto n \qquad \text{(constant P and T)}$$

Combining the three into one

Each law tells us how V depends on one variable with the others held fixed. If we now allow all three variables to change together, the combined dependence must be the product of the three separate proportionalities:

$$V \propto \frac{nT}{P}$$

To turn the proportionality into an equation we introduce a constant of proportionality, conventionally written as R, called the **universal molar gas constant** or simply the **gas constant.** Rearranging:

$$\mathbf{PV = nRT}$$

This is the **ideal gas equation**. It is also called the **ideal gas law** and the **equation of state** of an ideal gas, and the three names are used interchangeably. The word "state" here is a technical word: it means the condition of the gas as specified by its pressure, volume, and temperature together. An equation that links all of these in one relation deserves to be called the *equation of its state*.

The value of the constant R has been measured to great precision:

$$\mathbf{R = 8.314 J mol^{-1} K^{-1}}$$

The same value holds for every gas that behaves ideally. It is called the **universal** molar gas constant for precisely that reason, and the universality is a deep fact: it is a statistical consequence of the assumptions we laid down at the start of this chapter, and it would not be true if different gases had different molecular forces or different collision behaviours. Nitrogen, oxygen, helium, argon, and every other ideal gas obey $PV = nRT$ with the same R, which is why one equation serves them all.

For problems that compare two states of the same fixed amount of gas, it is often more convenient to write the equation in ratio form. Dividing $PV = nRT$ by nT gives $\frac{PV}{nT} = R$, which is the same constant for every state. Therefore, comparing state 1 with state 2:

$$\frac{P_1V_1}{n_1T_1} = \frac{P_2V_2}{n_2T_2}$$

If the amount of gas does not change between the two states (the gas is sealed), $n_1 = n_2$ and they cancel, leaving the **combined gas law:**

$$\mathbf{\frac{P_1V_1}{T_1} = \frac{P_2V_2}{T_2}}$$

The same equation in molecular units

The ideal gas equation as written above uses the number of moles n, which is a chemist's unit. For a physicist interested in the molecular picture, it is often more natural to use the total number of molecules N instead. The two are connected by Avogadro's constant:

$$\mathbf{N = n N_A} \text{ or } \mathbf{n = \frac{N}{N_A}}$$

Substituting $n = \frac{N}{N_A}$ into $PV = nRT$:

$$\mathbf{PV = \frac{N}{N_A} RT = NkT}$$

where:

$$\mathbf{k = \frac{R}{N_A} = \frac{8.314 \text{ J mol}^{-1} \text{ K}^{-1}}{6.022 \times 10^{23} \text{ mol}^{-1}} = 1.381 \times 10^{-23} JK^{-1}}$$

The constant is called the **Boltzmann constant**, named after the Austrian physicist Ludwig Boltzmann who laid most of the statistical foundations of the kinetic theory in the late nineteenth century. k is the gas constant expressed per molecule rather than per mole. Wherever a piece of physics involves the thermal energy of a single particle (atom, molecule, electron, photon), k is the constant that appears. It shows up in semiconductor physics, in the Boltzmann distribution of atomic energies, in the thermal noise voltage of a resistor, and in many other places far beyond gases. It is one of the most universal constants in physics, and it was born in this chapter, as a conversion factor between two different ways of counting the amount of gas.

The critical distinction between n and N

We now have two forms of the same equation:

$$PV = nRT \quad \text{(macroscopic form)}$$
$$PV = NkT \quad \text{(microscopic form)}$$

They describe the same gas and they give the same answer. The only difference is what unit of gas-amount they use. **n** counts **moles**; **N** counts **molecules**. A factor of the Avogadro constant (about 10^{23}) separates them, and this is where students confuse. Substituting n into a formula that wants N, or the other way round, gives an answer wrong by twenty-three orders of magnitude. A number like that is visibly absurd (a litre of gas does not contain 0.04 molecules, nor does it contain 2×10^{22} moles), so the error is usually caught, but not always, and the correction costs time in an examination.

Temperature must always be in Kelvin

Both forms of the equation require the **absolute** temperature, which means the temperature measured in kelvin, not in degrees Celsius or degrees Fahrenheit. The conversion is the one you already know:

$$T(K) = T(°C) + 273$$

This is the most common silent error in the whole chapter. A temperature of 27 °C is not 27 K; it is 300 K. A temperature of −3 °C is not a negative temperature at all in the ideal gas equation; it is 270 K. Substituting Celsius where Kelvin is wanted gives answers wrong by factors of ten or more, and unlike the n-versus-N error, this one is often silent: the answer is numerically plausible, just badly wrong. The only protection is habit. Every time you pick up a temperature, convert it to kelvin **first**, before you substitute, and write the kelvin value in your working so the examiner can see you did the conversion. If the question gives temperature in kelvin already (for example, "300 K"), write "T = 300K" again at the start of your working so you see that it is kelvin and not something else.

The quiet surprise: comparing the two equations

We now have two equations for the same quantity PV. On one hand, from the kinetic theory:

$$PV = \frac{1}{3}Nm\bar{c}^2$$

On the other hand, from the ideal gas equation in molecular form:

$$PV = NkT$$

Both are true for the same gas in the same state. They must agree. Setting the right-hand sides equal to each other:

$$\frac{1}{3}Nm\bar{c}^2 = NkT$$

The factor of N cancels from both sides, and multiplying by $3/2$:

$$\mathbf{\frac{1}{2}m\bar{c}^2 = \frac{3}{2}kT}$$

Look carefully at what this equation says. The left-hand side is the **average translational kinetic energy** of a **single** gas molecule (remember, $\bar{c}^2$ is the **mean** of the squares of the individual molecular speeds, so $\frac{1}{2}m\bar{c}^2$ is the **mean** kinetic energy of **a** molecule). The right-hand side is $\frac{3}{2}kT$, a quantity that depends on nothing but the temperature. It does not depend on the mass of the molecule, on the type of gas, on the pressure, or on the volume. Only on **T**.

This is the single most important result of the entire chapter. Temperature, on the molecular scale, is the average translational kinetic energy of the molecules, up to the constant factor $\frac{3}{2}k$. The thermometer you built in Chapter 2 was not measuring a mysterious thing called "hotness." *It was measuring, very indirectly, the average kinetic energy of the molecules hitting it*. **Hotness**, in the language of molecules, *means molecular speed.*

What we have just discovered is that the kinetic theory (which knows nothing about temperature) plus the ideal gas equation (which knows nothing about molecules) together *imply* that **temperature** *is directly related to molecular kinetic energy*. Neither piece says this on its own. Only the two, placed side by side, do. This is a derived result, and it is the payoff of the entire chapter so far. The next section unpacks the result in full and shows what it means for different kinds of gas, for different temperatures, and for the kinds of molecules you meet in everyday life.

Four examples follow.

BINDER Example 8

A sealed container of volume 0.050 m³ holds 2.0 mol of nitrogen gas (molar mass 28 g/mol) at a temperature of 300 K. Calculate: (a) the pressure of the gas, (b) the root mean square speed of the nitrogen molecules, and (c) verify numerically that the kinetic equation $PV = \frac{1}{3}Nm\bar{c}^2$ gives the same PV as the ideal gas equation.

Solution

(a) From $PV = nRT$:

$$P = \frac{nRT}{V} = \frac{(2.0\text{ mol})(8.314\text{ J mol}^{-1}\text{ K}^{-1})(300\text{ K})}{0.050\text{m}^3} = 9.98 \times 10^4\text{ Pa}$$

(b) The density of the gas (total mass divided by volume) uses the total mass nM:

$$\rho = \frac{nM}{V}$$

Substituting ρ to root mean square speed formula $v_{rms} = \sqrt{\frac{3P}{\rho}}$:

$$v_{rms} = \sqrt{\frac{3P}{\frac{nM}{V}}} = \sqrt{\frac{3PV}{nM}}$$

Substituting the numerical values:

$$v_{rms} = \sqrt{\frac{3 \times 9.98 \times 10^4 \text{Pa} \times 0.050 \text{m}^3}{2.0 \text{mol} \times 28 \times 10^{-3} \text{kg/mol}}} = 517 \text{m/s}$$

(c) For the verification step, let us compute $\frac{1}{3} Nm\bar{c}^2$ using $Nm = nM$ (the total mass of gas) and $\bar{c}^2 = v_{rms}^2$:

$$\frac{1}{3} Nm\bar{c}^2 = \frac{1}{3} (2.0 \times 28 \times 10^{-3} \text{kg})(517 \text{ m/s})^2 = 4990 \text{J}$$

Comparing with PV from the ideal gas equation:

$$PV = (9.98 \times 10^4 \text{Pa})(0.050 \text{m}^3) = 4990 \text{J}$$

The two expressions for PV agree, confirming that the kinetic equation and the ideal gas equation describe the same physical gas.

Making Sense of the Answer: *The pressure came out almost exactly one atmosphere, which is sensible for two moles of gas in fifty litres at room temperature (this is close to the molar volume at standard conditions). The rms speed of 517 m/s matches the canonical value for nitrogen at 300 K quoted in physics textbooks. The two independent calculations of PV agreeing numerically is not a coincidence; it is a check that the entire derivation of the previous section was arithmetically correct, and it is the kind of consistency check a careful student performs whenever they suspect a slip.*

Think Like a Physicist: *Whenever a chapter gives you two different-looking formulas for the same quantity, compute both and check they agree. This takes sixty seconds and catches arithmetic or conceptual errors immediately. The habit pays off enormously in long multi-part problems, where a silent error at step two can propagate through to the final answer without you noticing.*

BINDER Example 9

Two sealed containers, each of volume 0.040 m³, are held at the same temperature of 300 K. Container A holds 0.50 mol of helium (molar mass 4.0 g/mol). Container B holds 0.50 mol of argon (molar mass 40 g/mol). For each container, calculate (a) the pressure, (b) the root mean square speed of the atoms. Comment on which quantities are the same in the two containers and which are not, and why.

Solution

(a) Pressure from $PV = nRT$. Both containers have the same n, V, and T, so both have the same pressure:

$$P = \frac{nRT}{V} = \frac{(0.50 \text{ mol})(8.314 \text{ J mol}^{-1} \text{ K}^{-1})(300 \text{ K})}{0.040 \text{ m}^3} = 3.12 \times 10^4 \text{ Pa}$$

The pressure is the same in both containers, approximately 3.12×10^4 Pa.

(b) rms speed from $\mathbf{v_{rms}} = \sqrt{\frac{3PV}{nM}} = \sqrt{\frac{3nRT}{nM}} = \sqrt{\frac{\mathbf{3RT}}{\mathbf{M}}}$. Helium ($M = 4.0 \times 10^{-3}$ kg/mol):

$$v_{rms,He} = \sqrt{\frac{3 \times (8.314 \text{ J mol}^{-1} \text{ K}^{-1}) \times (300 \text{ K})}{4.0 \times 10^{-3} \text{ kg/mol}}} \approx 1368 \text{ m/s}$$

Argon ($M = 40 \times 10^{-3}$ kg/mol):

$$v_{rms,Ar} = \sqrt{\frac{3 \times (8.314 \text{ J mol}^{-1} \text{ K}^{-1}) \times (300 \text{ K})}{40 \times 10^{-3} \text{ kg/mol}}} \approx 433 \text{ m/s}$$

Equal and unequal quantities in the two containers:

- The **pressures** are *equal* because pressure depends only on n, V, T (from the ideal gas equation), and these are the same for both containers.
- The **average kinetic energies per atom** are *equal*, because average KE depends only on temperature ($\frac{3}{2}kT$), and the temperatures are identical.
- The **root mean square speeds** are *unequal*, because the atoms of helium are ten times lighter than the atoms of argon, and since $\frac{1}{2}mv_{rms}^2 = \frac{3}{2}kT$, a smaller mass forces a larger v_{rms}. Numerically, $\frac{v_{rms,He}}{v_{rms,Ar}} = \sqrt{\frac{40}{4}} = 3.16$, matching the ratio $\frac{1368}{433} = 3.16$ we computed.

Making Sense of the Answer: *The headline result is that two gases at the same temperature have the same average kinetic energy per molecule, but different molecular speeds, with the lighter gas moving faster. This is why helium balloons leak faster than air-filled balloons (the lighter helium atoms pass through microscopic pores in the rubber more quickly), and why hydrogen has escaped Earth's atmosphere over geological time while heavier oxygen has not.*

Think Like a Physicist: *When comparing two gases at the same temperature, the mass dependence of v_{rms} is the one thing that distinguishes them. Pressure depends only on n, V, T; kinetic energy per molecule depends only on T; the only species-dependent quantity is the rms speed, through the molar mass. If a question asks what is different between two gases at equal temperature, the answer almost always involves v_{rms} or some consequence of it (diffusion rate, effusion rate, speed of sound, escape from gravity).*

REAL Example 10

At the summit of Mount Kilimanjaro (altitude 5895m), the atmospheric pressure is approximately 48kPa, compared with 101kPa at sea level. The air temperature at the summit averages about 263K, compared with 301K at sea level in the surrounding Tanzanian lowlands. Mountaineers at the summit report significant difficulty breathing. Using the ideal gas equation in its molecular form, determine whether this difficulty is caused primarily by the lower number of air molecules per unit volume, or by the slower speeds of those molecules at the colder temperature. Justify your conclusion with an appropriate calculation.

Solution

The ideal gas equation in molecular form is $PV = NkT$, which rearranges to give the number density (molecules per unit volume):

$$\frac{N}{V} = \frac{P}{kT}$$

The ratio of number densities at the summit and at sea level:

$$\frac{\left(\frac{N}{V}\right)_{summit}}{\left(\frac{N}{V}\right)_{sea}} = \frac{\frac{P_{summit}}{kT_{summit}}}{\frac{P_{sea}}{kT_{sea}}} = \frac{P_{summit}}{P_{sea}} \times \frac{T_{sea}}{T_{summit}} = \frac{48kPa}{101kPa} \times \frac{301K}{263K} \approx 0.54$$

So the number density at the summit is about 54% of the sea-level value; there are 46% fewer air molecules per cubic metre. The temperature ratio alone (263/301 = 0.87) would only have reduced the number density by 13% if pressure had stayed the same. Most of the number-density drop comes from the lower pressure, not from the lower temperature.

For molecular speeds, $v_{rms} \propto \sqrt{T}$, so the ratio of rms speeds is:

$$\frac{v_{rms,summit}}{v_{rms,sea}} = \sqrt{\frac{T_{summit}}{T_{sea}}} = \sqrt{\frac{263K}{301K}} \approx 0.935$$

Air molecules at the summit move about 6.5% more slowly than at sea level, a modest change.

The difficulty of breathing at altitude is therefore primarily due to the lower **number density** of air molecules (46% fewer oxygen molecules available per breath) (and only minimally due to the slower motion of those molecules). A climber at the summit taking a normal breath draws about half as many oxygen molecules into the lungs as at sea level. The 6.5% slower molecular speeds are negligible in comparison.

Making Sense of the Answer: *The physics of altitude sickness is essentially a counting problem. Mountaineers struggle to breathe not because air molecules are cold and sluggish, but because there are far fewer of them per litre of lung volume. Supplemental oxygen tanks, used above 7000 m, work by restoring the number density of oxygen molecules reaching the lungs, not by warming them. The ideal gas equation in its molecular form, PV = NkT, turns what might seem a question about cold mountain air into a simple ratio of numbers of molecules per cubic metre.*

Think Like a Physicist: *Whenever pressure drops and temperature also drops, ask which is dropping faster. The combined gas law tells you that PV/T is constant, so a drop in P with a drop in T can partially cancel, leaving the number density almost unchanged, or it can compound, leaving the number density dramatically lower, depending on the relative sizes. The only way to tell which is by computing the ratio, as this example has done. Intuition is unreliable when two effects act in opposite directions on the same quantity.*

HOT Example 11

A sealed rigid gas cylinder in a storage yard holds gas at 20 °C and 12MPa. Kipute, walking past the yard one Saturday afternoon, notices that the cylinder has been standing in direct sunlight for several hours. She checks the gauge: the pressure has risen to 14MPa. The cylinder is rigid, so the volume of gas is unchanged, and it is sealed, so the amount of gas is unchanged. Find (a) the new temperature of the gas; (b) the average translational kinetic energy of a gas molecule at the original temperature of 20 °C; (c) the percentage increase in the root mean square speed of the molecules between the two states.

Solution

Convert the initial temperature to kelvin: $T_1 = 20 + 273 = 293$ K.

(a) At constant V and constant n, $PV = nRT$ reduces to $\frac{P}{T} = \text{constant}$:

$$\frac{P_1}{T_1} = \frac{P_2}{T_2}$$

$$T_2 = T_1 \times \frac{P_2}{P_1} = 293\,\text{K} \times \frac{14}{12} = 341.8\,\text{K} \approx 342\text{K} = 69\,^\circ\text{C}$$

The new temperature is about 69 °C.

(b) The average translational kinetic energy **per molecule** follows the molecular meaning of temperature established in this section, namely $\text{KE} = \frac{3}{2}\text{kT}$:

$$\text{KE} = \frac{3}{2} \times (1.38 \times 10^{-23}\,\text{J K}^{-1}) \times (293\,\text{K}) = 6.07 \times 10^{-21}\,\text{J}$$

Each molecule carries, on average, about 6×10^{-21} joules of kinetic energy.

This is a tiny amount individually; the energy of an entire mole of gas is obtained by multiplying by the Avogadro constant, giving roughly 3600J per mole.

(c) Since $v_{\text{rms}}^2 \propto T$, the rms speed satisfies $v_{\text{rms}} \propto \sqrt{T}$. The ratio of rms speeds is:

$$\frac{v_{\text{rms},2}}{v_{\text{rms},1}} = \sqrt{\frac{T_2}{T_1}} = \sqrt{\frac{342\text{K}}{293\text{K}}} = 1.080$$

The rms speed increases by a factor of 1.080, a percentage increase of 8.0%.

For comparison, the pressure has increased by the factor 14/12 = 1.167, a fractional increase of 16.7%, exactly matching the temperature ratio. Pressure scales linearly with Kelvin temperature; molecular speed scales only with its square root.

Making Sense of the Answer: *The cylinder has heated by 49K, and its pressure has jumped by 2MPa. This is why gas cylinders in Tanzania are required by safety regulations to be stored in shaded areas; a cylinder that was safe at 12MPa in the morning can exceed its pressure rating by noon if left in direct tropical sunlight. Kipute, noticing the problem, is being a thoughtful physicist. If the cylinder had been full to its rated pressure to begin with, the temperature rise could have caused a rupture, which is a serious industrial hazard.*

Think Like a Physicist: *Three quantities in this example scale with Kelvin temperature in three different ways. Pressure at constant V is linear in T; average kinetic energy per molecule is also linear in T; rms speed is linear in the square root of T. Whenever a problem asks you to compare two thermal states, stop and identify which of these three scalings applies to each quantity it asks about. The answer then writes itself as a ratio.*

Part (b) of the last example leaned on a result, $\text{KE} = \frac{3}{2}\text{kT}$, that we met only briefly at the end of the consistency argument. That result is rich enough to deserve a section of its own: it connects temperature to molecular speeds, to the speeds of different gases at the same temperature, to escape velocities and planetary atmospheres, and to the eventual definition of internal energy for an ideal gas. The next section is devoted to unpacking it.

TEMPERATURE AND MOLECULAR KINETIC ENERGY

Chapter 2 asked what temperature is and gave an answer that was, on reflection, a little evasive. Temperature was defined operationally, as whatever a thermometer measures, with the zero point of the Celsius scale pegged to the freezing of water and the zero point of the Kelvin scale pegged to the absolute zero of thermal motion. That was enough to do thermometry, but it did not say what temperature actually is. A thermometer reading of 300K was a number on a scale; it was not an explanation. This section removes the evasion. We have now assembled enough machinery to say what temperature really is on the molecular scale, and the answer, when it comes, is disarmingly simple. **Temperature** *is molecular translational kinetic energy, multiplied by a constant.* That is the whole of it.

The result we are about to unpack was quietly stated at the end of the previous section, when we equated the kinetic equation with the ideal gas equation. We now take it seriously, develop its immediate consequences, and use it to answer a question that has been hanging over the chapter since the first paragraph: why does Kipanga's pump get hot when he compresses the air, and what does "hot" actually mean down at the level of the molecules themselves?

The key result

From the two equations for PV:

$$\frac{1}{3}\,\text{Nm}\bar{c}^2 = \text{NkT}$$

Cancelling N and multiplying by $3/2$:

$$\frac{1}{2}\,\text{m}\bar{c}^2 = \frac{3}{2}\text{kT}$$

The left side is the average translational kinetic energy of a single molecule of the gas, which we will denote KE. The right side is three-halves of the Boltzmann constant times the Kelvin temperature. So:

$$\text{KE} = \frac{3}{2}\text{kT} \left(\text{or T} = \frac{2}{3\text{k}} \times \text{KE} = \text{constant} \times \text{KE}\right)$$

This is the central result of the chapter. Read it slowly, because it says more than it looks like it says.

What this equation is saying

The average translational kinetic energy per molecule of an ideal gas depends on one thing and one thing only: the **Kelvin** temperature. It does not depend on the mass of the molecule. It does not depend on the chemical identity of the gas. It does not depend on the pressure. It does not depend on the volume. It does not depend on whether the gas is in a balloon, a cylinder, or a tyre. Two different gases at the **same** temperature, whether they are helium and argon, or nitrogen and xenon, or hydrogen and mercury vapour, have **exactly** the same average translational kinetic energy per molecule. There is no small correction. There is no qualifying clause. Temperature is the measure of translational kinetic energy per molecule, and nothing more.

This is a surprising claim when you first meet it, because it seems to conflict with everyday intuition. Heavier molecules should surely carry more energy at a given temperature, you might think. They don't. What does change with molecular mass is the **speed** of the molecules: *lighter molecules move faster, heavier molecules move more slowly, in just the right proportion to keep the kinetic energy* $\frac{1}{2}m\bar{c}^2$ *the same for both.* A hydrogen molecule and an oxygen molecule at room temperature are both carrying about 6×10^{-21}J of kinetic energy on average. The hydrogen molecule, being sixteen times lighter, moves four times faster to make up the difference. This is the numerical fact we will confirm with the second worked example of this section.

Three formulas for the root mean square speed

Rearranging the central result for the rms speed of a single molecule $\left(\frac{1}{2}\text{m}\bar{c}^2 = \frac{3}{2}\text{kT}\right)$:

$$\mathbf{v_{rms}} = \sqrt{\bar{\mathbf{c}}^{\mathbf{2}}} = \sqrt{\frac{\mathbf{3kT}}{\mathbf{m}}}$$

This is the first form, and it expresses the rms speed in terms of the per-molecule mass m and the Boltzmann constant k. For numerical calculations it is sometimes more convenient to use the molar mass M instead of the per-molecule mass, using $\text{m} = \frac{\text{M}}{\text{N}_\text{A}}$ and $\text{k} = \frac{\text{R}}{\text{N}_\text{A}}$:

$$\mathbf{v_{rms}} = \sqrt{\frac{3 \times \frac{\text{R}}{\text{N}_\text{A}} \times \text{T}}{\frac{\text{M}}{\text{N}_\text{A}}}} = \sqrt{\frac{\mathbf{3RT}}{\mathbf{M}}}$$

This second form is the most practical for calculations, because molar masses are tabulated in every physics and chemistry reference while per-molecule masses require a division by N_A that adds work without insight. Finally, from the pressure derivation we already have the density form:

$$\mathbf{v_{rms}} = \sqrt{\frac{\mathbf{3P}}{\boldsymbol{\rho}}}$$

Three equivalent formulas, each matched to a different question. Use the first when you know the per-molecule mass (rare). Use the second when you know the molar mass and the temperature (the usual NECTA case). Use the third when you know the pressure and the density (the usual lab case).

Total kinetic energy of a gas sample

Average kinetic energy per molecule is $\frac{3}{2}$kT. The total translational kinetic energy of a gas sample containing N molecules is obtained by multiplying:

$$\mathbf{KE_{total}} = \text{N} \times \frac{3}{2}\text{kT} = \mathbf{\frac{3}{2}\,NkT} = \mathbf{\frac{3}{2}\,nRT}$$

where the last equality uses $\text{NkT} = \text{nRT}$. For one mole of gas at room temperature ($\text{T} = 298\text{K}$), the total translational kinetic energy is $\frac{3}{2}(1)(8.314)(298) \approx 3\,716$J, or about 3.7kJ. This is the same kind of energy as the kinetic energy of a one-kilogram mass moving at $\sqrt{2 \times 3716} \approx 86\text{m/s}$. A litre of air at room temperature carries, in its molecular motion, roughly the kinetic energy of a hard-thrown cricket ball. This is why compressing a gas heats it noticeably: that molecular energy is genuinely substantial, and mechanical work done on the gas adds to it.

With this central result now firmly in hand, let us bring it down to earth through a few carefully chosen worked examples.

BINDER Example 12

Calculate the average translational kinetic energy of a gas molecule at (a) 0 °C and (b) 100 °C. Find the ratio of the two energies and explain why it is not 2, despite the temperature having apparently doubled between freezing and boiling water.

Solution

Converting to kelvin: $T_1 = 273\text{K}$, $T_2 = 373\text{K}$.

(a) At 0 °C:

$$KE_1 = \frac{3}{2}kT = \frac{3}{2} \times (1.381 \times 10^{-23}\,\text{JK}^{-1}) \times (273\text{K}) = 5.65 \times 10^{-21}\text{J}$$

(b) At 100 °C:

$$KE_2 = \frac{3}{2} \times (1.381 \times 10^{-23}\,\text{J K}^{-1}) \times (373\text{K}) = 7.73 \times 10^{-21}\text{J}$$

Ratio:

$$\frac{KE_2}{KE_1} = \frac{373}{273} = 1.37$$

The kinetic energy at 100 °C is 1.37 times the kinetic energy at 0 °C, not 2 times, because the Celsius temperature of 100 °C is not twice 0 °C on the absolute scale. On the Kelvin scale, which is the one the kinetic theory uses, 0 °C is 273 K and 100 °C is 373 K, and 373 is only 1.37 times 273. The Celsius scale, because its zero point is arbitrary, gives misleading ratios. This is why kinetic theory always uses Kelvin.

Making Sense of the Answer: *The 37% increase in kinetic energy between freezing water and boiling water is modest; it is not the factor of two that a naive reading of "0 to 100" suggests. This is why heating a pot of water from 0 °C to 100 °C takes far less energy than heating it from 0 K to 273 K. Most of the thermal energy of a substance at everyday temperatures was already there at the freezing point.*

Think Like a Physicist: *Ratios of Celsius temperatures are physically meaningless. Ratios of Kelvin temperatures are the real thing. Whenever a question asks by what factor some thermal quantity has changed, first rewrite the temperatures in kelvin, then take the ratio.*

BINDER Example 13

Hydrogen gas (molar mass 2.0 g/mol) and oxygen gas (molar mass 32 g/mol) are both held at a temperature of 300 K. Calculate the root mean square speed of each gas. Then calculate the average translational kinetic energy per molecule for each gas, and verify numerically that the two values are equal.

Solution

rms speeds from $v_{rms} = \sqrt{\frac{3RT}{M}}$. For hydrogen ($M = 2.0 \times 10^{-3}$kg/mol):

$$v_{rms,H_2} = \sqrt{\frac{3 \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K})}{2.0 \times 10^{-3}\text{kg/mol}}} = 1934\text{m/s}$$

For oxygen ($M = 32 \times 10^{-3}$kg/mol):

$$v_{rms,O_2} = \sqrt{\frac{3 \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K})}{32 \times 10^{-3}\text{kg/mol}}} = 484\text{m/s}$$

Hydrogen molecules move four times faster than oxygen molecules (ratio 1934/484 ≈ 4.00, matching $\sqrt{\frac{32}{2}}$ since oxygen is sixteen times heavier per molecule).

Kinetic energy per molecule uses $\frac{1}{2}mv_{rms}^2$.

Per-molecule masses:

$$m_{H_2} = \frac{M}{N_A} = \frac{2.0 \times 10^{-3}\text{kgmol}^{-1}}{6.022 \times 10^{23}\text{mol}^{-1}} = 3.32 \times 10^{-27}\text{kg}, m_{O_2} = \frac{32 \times 10^{-3}\text{kgmol}^{-1}}{6.022 \times 10^{23}\text{mol}^{-1}} = 5.31 \times 10^{-26}\,\text{kg}$$

$$KE_{H_2} = \frac{1}{2} \times (3.32 \times 10^{-27}\text{kg}) \times (1934\text{m/s})^2 = 6.21 \times 10^{-21}\text{J}$$

$$KE_{O_2} = \frac{1}{2} \times (5.31 \times 10^{-26}\text{kg}) \times (484\text{m/s})^2 = 6.22 \times 10^{-21}\text{J}$$

The two kinetic energies are equal to within rounding, as the theory predicts.

Understand temperature confirmation: Both can be verified independently against the temperature formula:

$$KE = \frac{3}{2}kT = \frac{3}{2} \times (1.381 \times 10^{-23}\,\text{J K}^{-1}) \times (300\text{K}) = 6.21 \times 10^{-21}\text{J}$$

All three numbers agree. The critical insight is confirmed: *at a given temperature, different gases have different molecular speeds, but the same kinetic energy per molecule.*

Making Sense of the Answer: *Hydrogen molecules at room temperature move at nearly two kilometres per second, which is comparable to a rifle bullet at the moment it leaves the muzzle. Oxygen molecules, being sixteen times heavier, move at only a quarter of that speed, but with exactly the same kinetic energy per molecule. The speed difference is what matters for diffusion rates and for escape from gravity; the kinetic energy equality is what matters for thermal equilibrium. Both will come back in the next example.*

Think Like a Physicist: *When two gases are in thermal equilibrium, ask what is equal at the molecular level and what is not. Temperature is equal, therefore kinetic energy per molecule is equal. Masses are not equal, therefore speeds are not equal. Number of molecules is set by partial pressures, not by temperature. Identifying these three separately untangles most kinetic-theory problems on sight.*

REAL Example 14

Earth's atmosphere contains significant amounts of oxygen and nitrogen but only trace amounts of hydrogen and helium, even though hydrogen and helium are far more abundant in the Universe as a whole. The escape velocity from Earth's gravity is approximately 11.2km/s. Using the result of the previous example, explain why hydrogen can escape Earth's atmosphere over geological time while oxygen cannot.

Solution

The root mean square speed of hydrogen molecules at room temperature is about 1.9 km/s (from the previous example), which is approximately 17% of Earth's escape velocity of 11.2 km/s. The corresponding figure for oxygen is 0.48 km/s, only about 4% of escape velocity. Although neither rms speed exceeds escape velocity, the rms is not the full story: molecular speeds in a gas follow a distribution, with some molecules moving slower than the rms speed and some moving faster. In the high-speed tail of the distribution, a small fraction of molecules always exceeds any given speed, including the escape velocity. For hydrogen, whose rms is already 17% of escape, this fraction is non-negligible: over geological time (billions of years), enough hydrogen molecules in the high-altitude tail of the distribution exceed escape velocity, travel upward, and leave the atmosphere permanently. For oxygen, whose rms is only 4% of escape, the fraction of the distribution exceeding escape velocity is astronomically smaller, and the rate of escape is effectively zero on any timescale relevant to a planet.

It is worth understanding that: The same reasoning explains why the Moon has almost no atmosphere (its escape velocity is only 2.4 km/s, which even oxygen can exceed over geological time), why Jupiter retains even hydrogen (its escape velocity is 59.5 km/s, far above any molecular speed at its cold outer temperatures), and why the top of Mount Kilimanjaro, although thinner in number density, has the same chemical composition as the lowlands: 5895m is not nearly high enough for the escape-velocity-plus-distribution argument to matter, because molecules escape only from the very top of the atmosphere, far above any mountain.

Making Sense of the Answer: *The chemical composition of a planet's atmosphere is set by a competition between thermal kinetic energy and gravitational binding energy. Lighter molecules, which move faster at the same temperature, lose this competition over time. Earth has retained enough oxygen and nitrogen to support life largely because these molecules are heavy enough to stay gravitationally bound. Had Earth been smaller, or warmer, or both, we would not be here to have this conversation.*

Think Like a Physicist: *When an average quantity is smaller than some threshold, the tail of the distribution often matters more than the average itself. Averages describe typical molecules; tails describe rare events that can be the only thing that matters over long timescales. Atmospheric escape is a rare-event phenomenon, and so is radioactive decay, and so is thermal fluctuation in electronic noise, and so are many other phenomena in physics where the average is large but the tail is what the universe notices.*

HOT Example 15

At what temperature must nitrogen gas (molar mass 28g/mol) be heated in order for its root mean square molecular speed to equal that of hydrogen gas at 300K? Comment briefly on the size of your answer.

Solution

Using:

$$v_{rms} = \sqrt{\frac{3RT}{M}}$$

The condition we are asked to satisfy is equality of root mean square speeds between the two gases at their respective temperatures:

$$v_{rms,N_2} \text{ at } T = v_{rms,H_2} \text{ at } 300K$$

Squaring both sides (so we can use $v_{rms}^2 = \frac{3RT}{M}$ directly) and writing out each side:

$$\frac{3R \times T}{M_{N_2}} = \frac{3R \times 300K}{M_{H_2}}$$

The factor 3R cancels from both sides, leaving a clean ratio equation. Solving for T:

$$T = \frac{M_{N_2}}{M_{H_2}} \times 300K = \frac{28 gmol^{-1}}{2.0 gmol^{-1}} \times 300K = 14 \times 300K = 4200K$$

The temperature required is about fourteen times room temperature and comparable to the surface temperature of the Sun (which is about 5800K). This shows that heavy molecules would need extremely high temperatures to match the speed of light molecules at ordinary temperatures, which is why light gases such as hydrogen escape from planetary atmospheres while heavier ones are retained.

Making Sense of the Answer: *the temperature ratio 4200/300 = 14 is exactly the molar mass ratio 28/2 = 14, as it must be from $v_{rms} \propto T/M$. This is the same scaling law that governs atmospheric escape, stellar composition, and the speed of sound in different gases. A single quadratic proportionality, $\bar{c}^2 \propto T/M$, carries an enormous amount of astrophysics inside it.*

Think Like a Physicist: *whenever two molecular speeds are equal, the product MT must be equal. This "equal-speed" condition often appears disguised in problems about diffusion, effusion, or spectral broadening. Recognising that MT is the invariant, not M or T separately, turns a one-line calculation into a sentence.*

Having enjoyed the worked examples, we now realise that something has been quietly left unsaid. Every result in this section has been about translational motion. A molecule that only translates, like a helium atom, is fully described by the story we have just told. But many gases are made of molecules that can also rotate, like the dumbbell-shaped nitrogen and oxygen that fill the air around you, and some can even vibrate internally, like the bent-shape water vapour and the linear carbon dioxide. These additional ways of storing energy do not show up in the pressure formula, but they do show up in the total energy a gas carries at a given temperature. The next subtopic makes the correction, and in doing so introduces an integer that will quietly control a great deal of physics still to come.

DEGREES OF FREEDOM AND INTERNAL ENERGY

Every result so far in this chapter has been about translational motion. The kinetic equation counted the momentum transferred to a wall when a molecule bounced off it, and that bouncing is translational motion. The central result $KE = \frac{3}{2} kT$ is the average **translational** kinetic energy per molecule, not the total energy. For a **monatomic** gas like helium or argon, whose atoms are essentially spherical point-masses with nothing inside them to wobble or spin, this is the whole story. A helium atom at 300 K carries $\frac{3}{2}$ kT of energy and nothing more.

But most real gases are not monatomic. Nitrogen and oxygen, the two gases that make up 99% of the air you are breathing, are **diatomic**: each molecule is a pair of atoms held together by a chemical bond, something like a tiny dumbbell. A dumbbell can translate (the whole molecule moves through space), but it can also **rotate** end over end, and the energy stored in that rotation does not show up in the pressure formula. Carbon dioxide, water vapour, methane, and ozone are **polyatomic:** three or more atoms in a single molecule, with even more ways to move internally. This section extends the kinetic theory to account for those extra motions, introduces the concept of **degrees of freedom** as a bookkeeping tool, and arrives at a compact formula for the total internal energy of any ideal gas in terms of a single integer.

What a degree of freedom is

A **degree of freedom** of a molecule *is an independent way in which the molecule can store kinetic energy*. For a point-like atom moving through three-dimensional space, there are three such ways: motion along the x-axis, motion along the y-axis, and motion along the z-axis. Each of these is independent, in the sense that the energy associated with motion along one axis does not appear in any other. The total translational kinetic energy of the atom is a sum of three independent contributions, one per axis.

Translational motion gives three degrees of freedom. Rotational motion adds more, provided the molecule can actually rotate in a distinguishable way. A single spherical atom has no distinguishable axis of rotation (a perfect sphere looks the same when rotated about any axis through its centre), so a monatomic molecule has zero rotational degrees of freedom. A **diatomic molecule**, shaped like a dumbbell, *has **two** perpendicular axes about which rotation is distinguishable* (both perpendicular to the bond axis); rotation about the bond axis itself stores negligible energy because the atoms are essentially points on the bond, so only two rotational degrees of freedom contribute. A **polyatomic molecule** *has **three** rotational degrees of freedom*, because its non-linear arrangement of atoms distinguishes three perpendicular axes of rotation.

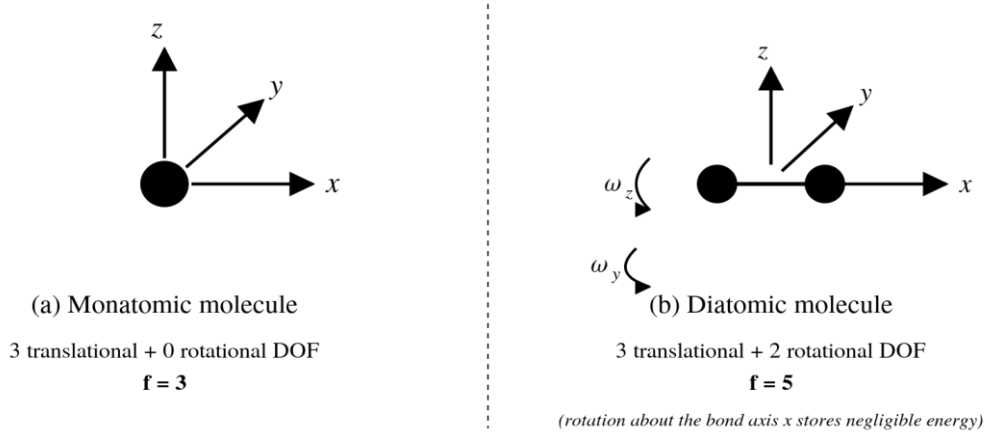


Figure: ***Degrees of freedom of a monatomic molecule (left) and a diatomic molecule (right).*** *The monatomic molecule has only three translational degrees of freedom. The diatomic molecule has three translational and two rotational degrees of freedom (about axes perpendicular to the bond).*

The equipartition theorem, stated informally

Having identified the degrees of freedom, the question is how much energy each one stores. The answer is a deep result of statistical mechanics, called the **equipartition theorem**, which we state here without proof:

In thermal equilibrium, each degree of freedom of a molecule stores, on average, $\frac{1}{2}\mathrm{kT}$ *of energy per molecule.*

This statement is consistent with what we already know. The three translational degrees of freedom give an average translational kinetic energy of $3 \times \frac{1}{2}\mathrm{kT} = \frac{3}{2}\mathrm{kT}$ per molecule, exactly the result from the previous section. The equipartition theorem generalises this: whatever other degrees of freedom a molecule has, each one also contributes $\frac{1}{2}$ kT. A full proof of the theorem requires the Boltzmann distribution of energies in a gas and the equipartition of energy among independent quadratic terms in the Hamiltonian, and lies beyond this chapter; the theorem is revisited in Chapter 5 in connection with heat capacities. For now, we take it as given and use it.

Total kinetic energy per molecule

For a molecule with a total number of degrees of freedom $\mathrm{f} = \mathrm{f_{trans}} + \mathrm{f_{rot}}$, the total average kinetic energy per molecule is:

$$\mathrm{KE_{tot}} = \frac{\mathrm{f}}{2}\mathrm{kT}$$

Filling in the three gas types:

- Monatomic gases (helium, argon, neon) have $\mathrm{f} = 3 + 0 = 3$, so $\mathrm{KE_{tot}} = \frac{3}{2}\mathrm{kT}$.
- Diatomic gases (hydrogen, nitrogen, oxygen) have $\mathrm{f} = 3 + 2 = 5$, so $\mathrm{KE_{tot}} = \frac{5}{2}\mathrm{kT}$.
- Polyatomic gases (carbon dioxide, water vapour, ozone) have $\mathrm{f} = 3 + 3 = 6$, so $\mathrm{KE_{tot}} = 3\mathrm{kT}$.

Internal energy of a gas sample

The **internal energy** U of a gas sample *is the total kinetic energy of all its molecules.* For N molecules each carrying $\frac{\mathrm{f}}{2}\mathrm{kT}$:

$$\mathbf{U} = \mathrm{N} \times \frac{\mathrm{f}}{2}\,\mathrm{kT} = \mathbf{\frac{f}{2}NkT} = \mathbf{\frac{f}{2}nRT}$$

Three compact results for the three gas types follow immediately. For a monatomic gas ($\mathrm{f} = 3$):

$$\mathrm{U} = \frac{3}{2}\mathrm{nRT} \qquad \text{(monatomic)}$$

For a diatomic gas ($\mathrm{f} = 5$):

$$\mathrm{U} = \frac{5}{2}\mathrm{nRT} \qquad \text{(diatomic)}$$

For a polyatomic gas ($\mathrm{f} = 6$):

$$\mathrm{U} = 3\mathrm{nRT} \qquad \text{(polyatomic)}$$

Three formulas, controlled by a single integer f. That integer is carrying enormous predictive power, as we will see at the end of this section and in Chapter 5.

Summary of internal energy for the three gas types

Gas type	Examples	Degrees of freedom *f*	Internal energy per mole
Monatomic	He, Ne, Ar	3 (3 trans + 0 rot)	U = (3/2) RT
Diatomic	H_2, N_2, O_2	5 (3 trans + 2 rot)	U = (5/2) RT
Polyatomic	CO_2, H_2O, O_3	6 (3 trans + 3 rot)	U = 3 RT

With these ideas now neatly in place, let us anchor them through a few worked examples before they drift away.

BINDER Example 16

Calculate the total internal energy of 3.0 mol of helium gas (monatomic) at a temperature of 27 °C.

Solution

$T = 27 + 273 = 300K$. Helium is monatomic, so $f = 3$ and the internal energy is:

$$U = \frac{3}{2}nRT = \frac{3}{2} \times (3.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 11224J$$

The total internal energy of 3.0 mol of helium at 300 K is approximately 1.12×10^4 J (or about 11.2 kJ).

Making Sense of the Answer: *Eleven kilojoules is the energy that would be released if the three moles of helium gas were to be cooled all the way to absolute zero. It is comparable to the chemical energy in about 0.7 grams of sugar, which is a small but non-trivial quantity. Gas at room temperature is carrying real energy, distributed across an astronomical number of molecules in kinetic form.*

Think Like a Physicist: *The internal energy of a monatomic gas depends only on n, T, and R, with no dependence on the gas identity or molar mass. Helium, neon, argon, and xenon in equal numbers of moles at the same temperature all have exactly the same internal energy. This is the cleanest possible gas.*

BINDER Example 17

Calculate the total internal energy of 2.0 mol of nitrogen gas (diatomic) at 27 °C. By what factor is this larger than the internal energy of 2.0 mol of helium gas at the same temperature?

Solution

Convert temperature: $T = 300K$. Nitrogen is diatomic, so $f = 5$ and:

$$U_{N_2} = \frac{5}{2}nRT = \frac{5}{2} \times (2.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 12471J$$

For 2.0 mol of helium at the same temperature:

$$U_{He} = \frac{3}{2}nRT = \frac{3}{2} \times (2.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 7483J$$

Ratio of internal energies:

$$\frac{U_{N_2}}{U_{He}} = \frac{12471}{7483} = \frac{5}{3} \approx 1.67$$

The nitrogen has 5/3 times the internal energy of the helium at the same temperature, purely because each nitrogen molecule has two additional rotational degrees of freedom that helium atoms do not have. The ratio 5/3 is exactly the ratio of the degrees of freedom.

Making Sense of the Answer: *Two gases at the same temperature and in the same amount have different internal energies if they have different molecular structures. A diatomic gas stores 67% more energy per mole than a monatomic gas, because its molecules have extra rotational motion that the monatomic atoms cannot participate in. This is why heating a given volume of nitrogen from 0 °C to 100 °C requires more heat than heating the same volume of helium over the same temperature range. The rotational motion is a storage reservoir that the heat has to fill along with the translational one.*

Think Like a Physicist: *The ratio $U_{diatomic}/U_{monatomic}$ at equal n and T equals the ratio of degrees of freedom, 5/3. The absolute values require a full substitution; the ratio requires no numerical work at all. When two gas samples differ only in molecular structure at matching n and T, reach for the ratio first.*

REAL Example 18

Account for the following:

(a) It takes longer to heat a room full of humid air than dry air to the same temperature rise.
(b) A helium-filled balloon released indoors warms up faster than an air-filled balloon of the same size as the room temperature changes.
(c) Monatomic gases such as argon are used as inert fillers in insulating double-glazed windows rather than air, despite being more expensive.

Solution

(a) Dry air is mostly diatomic (nitrogen and oxygen, f = 5), while water vapour is polyatomic (f = 6) and stores more internal energy per mole per kelvin than the diatomic molecules. Humid air therefore requires more heat than dry air to undergo the same temperature rise, so heating a humid room takes longer.
(b) Helium is monatomic (f = 3) while air is effectively diatomic (f = 5), so helium needs less heat per mole to raise its temperature by a given amount. For the same heat flow through the balloon wall, the helium-filled balloon therefore reaches the room's temperature faster than the air-filled balloon.
(c) Argon atoms are monatomic (f = 3) and store less thermal energy per collision than the diatomic nitrogen and oxygen molecules of air (f = 5). In addition, argon atoms are heavier than the average air molecule, so at the same temperature they move more slowly and transport heat across the gap less rapidly. Both effects lower the thermal conductivity of argon below that of air, making argon a better insulator despite its greater cost.

Making Sense of the Answer: *The three situations look superficially different, but they are all controlled by the same integer f. Humid air stores more energy per kelvin because water vapour has more degrees of freedom. Helium balloons warm up faster because helium has fewer degrees of freedom. Argon is a better insulator because its monatomic structure carries less heat per molecule. The physical question in each case is how many independent ways a molecule can absorb, store, or transport thermal energy, and f is the answer.*

Think Like a Physicist: *Whenever a real-world thermal question involves "this substance heats faster than that one" or "this gas insulates better than that one", ask first which molecules have more degrees of freedom and which are heavier. The answer is usually a combination of the two, and reasoning about both factors separately is clearer than reaching for a table of heat capacities. Molecular structure, through f and m, sits underneath most everyday thermal phenomena even when the phenomena themselves look unrelated.*

HOT Example 19

A rigid sealed container of volume 0.040 m³ holds a mixture of 1.0 mol of helium gas and 1.0 mol of nitrogen gas, both at a temperature of 300 K. Each gas behaves ideally and neither reacts with the other. Calculate:

(a) the total internal energy of the mixture,

(b) the total pressure exerted on the walls of the container.

Solution

(a) The total internal energy is the sum of the internal energies of the two component gases, because the two gases do not interact. Helium is monatomic ($f = 3$); nitrogen is diatomic ($f = 5$):

$$U_{tot} = U_{He} + U_{N_2} = \frac{3}{2} n_{He} RT + \frac{5}{2} n_{N_2} RT$$

Substituting numerical values:

$$U_{tot} = ½ \times (3 \times 1.0\text{mol} + 5 \times 1.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K}) = 9977\text{J}$$

The total internal energy of the mixture is approximately 9977J.

(b) For non-reacting helium-nitrogen gaseous mixture:

$$P_{tot} = \frac{(n_{He} + n_{N_2}) RT}{V}$$

$$P_{tot} = \frac{(2.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K})}{0.040\text{m}^3} = 1.25 \times 10^5\text{Pa}$$

The total pressure on the walls of the container is approximately $1.25 \times 10^5\text{Pa}$.

Making Sense of the Answer: *The total pressure of the mixture is the same as it would be for 2.0 mol of a single ideal gas in the same container at the same temperature. Pressure depends on n, V, and T alone; it does not care whether the molecules are all the same kind or a mixture of different kinds. Internal energy, by contrast, does care: because helium and nitrogen have different numbers of degrees of freedom, they carry different amounts of internal energy per mole even at the same temperature. Pressure is blind to molecular structure; internal energy is not.*

Think Like a Physicist: *For any gas mixture at uniform temperature, each component obeys PV = nRT independently, with its own partial pressure contributed to the total. This is Dalton's law of partial pressures, and it turns multi-component gas problems into separate single-component problems that are added at the end. The algebra never gets harder as more gases are added; it only gets longer.*

Everything in this section has depended on a single integer, the number of degrees of freedom f. Monatomic gases have f = 3; diatomic gases have f = 5; polyatomic gases have f = 6. The internal energy of any ideal gas is then (f/2)nRT, and that is the end of the story at this level. But the same integer f will reappear in Chapter 5 in the formulas for the heat capacity at constant volume and the heat capacity at constant pressure of an ideal gas. The humble integer f, counting the ways a molecule can store kinetic energy, turns out to control also how much heat a gas absorbs when warmed, and how fast sound travels through it, and how adiabatic compression changes its temperature. Carrying f correctly through this section is not just a bookkeeping habit; it is preparation for a chapter's worth of physics that is yet to come.

With the ideal gas story now fully told, we find ourselves at a natural moment to ask a harder question. We have spent five sections assuming that molecules have no volume of their own, that they do not attract each other, and that all collisions are perfectly elastic. These assumptions produced an equation of state so clean and predictive that it is hard to remember they were assumptions at all. But every real gas departs from this idealisation under the right conditions, and the departures are neither small nor exotic. The next subtopic asks where the ideal gas model breaks down, why it breaks down, and what the real physics looks like beyond the approximation.

REAL GASES AND THE LIMITS OF THE IDEAL GAS MODEL

Everything we have done so far has rested on the six assumptions of the kinetic theory. Those assumptions gave us a theory clean enough to derive, predictive enough to use, and simple enough to remember. But no real gas obeys them perfectly. A real gas is made of molecules that do occupy a small volume of their own, that do attract each other weakly even when they are not colliding, and that do, at sufficiently low temperatures, stop moving in straight lines and instead form a liquid. The ideal gas equation is a beautiful, but it is an approximation, and this section looks honestly at where the approximation breaks down and why.

Which assumptions fail first

When a real gas is pushed away from ideal behaviour, it is not all six assumptions that fail at once. Two of them fail first, and the other four remain adequate until very extreme conditions are reached.

The first to fail is **assumption 2**, that the volume of the molecules themselves is negligible compared with the volume of the container. At ordinary pressures this is excellent: one mole of gas at atmospheric pressure and room temperature occupies about 25 litres, while the actual molecules inside it occupy only a few cubic centimetres. But at **high pressure** *the container shrinks around the gas while the molecules themselves do not shrink, and the fraction of the container occupied by the molecules becomes significant.* A gas compressed to one hundred atmospheres has molecules occupying perhaps a few percent of the container volume, and at one thousand atmospheres the fraction is larger still. The molecules begin to get in each other's way, and the equation $PV = nRT$ no longer holds to the precision the assumption claimed.

The second to fail is **assumption 4**, that intermolecular forces are negligible except during the instant of a collision. This assumption fails at **low temperature.** At room temperature, molecules are moving fast enough that the weak attractive forces between them (van der Waals forces, named after the physicist who first wrote them into an equation of state) are too feeble to deflect a molecule noticeably from its straight-line path. But at **low temperatures** *the molecules are moving slowly, and the attractive forces have time to act.* Molecules begin to stick to each other briefly, to cluster, and at sufficiently low temperatures to condense into a liquid altogether. The gas stops behaving as a collection of independent projectiles and starts behaving as a collection of mutually attracting particles.

The other four assumptions (large number of molecules, elastic wall collisions, straight-line motion between collisions, negligible collision time) remain adequate well into the regime where (2) and (4) have failed, and they are not usually what limits the ideal gas model in practice.

Conditions for near-ideal behaviour

The two failure modes just described point directly at the conditions under which a real gas behaves most like an ideal gas. The ideal gas approximation is best at **low pressure**, because *low pressure means a large container volume for a given amount of gas, so the molecules occupy a negligible fraction of that volume.* It is best at **high temperature**, because *at high temperature the molecules move fast and the weak intermolecular attractions have no time to act on them.* It is best at **low density**, because *low density combines both advantages: molecules are far apart (so their own volume is negligible) and typically moving fast (so attractions do not matter).*

These three conditions are equivalent statements of the same physical idea, because low pressure at a given temperature implies low density, and high temperature at a given density implies faster-moving molecules less affected by attractions.

Students often remember the compact phrase: *real gases behave most ideally at low pressure and high temperature*. That is enough for the examination.

The van der Waals equation

The Dutch physicist Johannes van der Waals proposed in 1873 a modified equation of state that corrects for the two effects identified above. Schematically, his equation takes the form:

$$(\mathrm{P} + \text{correction for attractions}) \times (\mathrm{V} - \text{correction for molecular volume}) = \mathrm{nRT}$$

The first correction raises the effective pressure to account for attractive forces between molecules; the second reduces the available volume to account for the space the molecules themselves take up. At low pressure and high temperature, both corrections become small and the equation reduces back to $\mathrm{PV} = \mathrm{nRT}$. The detailed form of the corrections, and their derivation, lie beyond the A-Level syllabus and are not required for the examination. What matters here is the conceptual point: the van der Waals equation exists, it works better than $\mathrm{PV} = \mathrm{nRT}$ for real gases away from the ideal regime, and it does so by restoring the two assumptions the ideal gas model simplified away.

An ideal gas cannot be liquefied

One further consequence of the failure of assumption (4) deserves its own paragraph. A real gas can be liquefied by cooling it sufficiently and, if necessary, compressing it; every industrial gas sold in Tanzania has a liquid phase, produced in this way for **storage** and **transport**. An ideal gas, by contrast, can never be liquefied. This follows directly from assumption (4): if intermolecular forces are truly zero, as the ideal gas model requires, then there is nothing to hold the molecules together into a liquid. No matter how cold or how compressed the ideal gas becomes, its molecules continue to bounce off each other elastically and move in straight lines between collisions. They never stick. The ability of real gases to condense into liquids is itself the clearest evidence that intermolecular forces are not zero, and therefore that no real gas is strictly ideal.

With the honest limits of the ideal gas model now acknowledged, we can return to the many situations where the model is perfectly adequate and see what it can do. The kinetic theory of gases is not just an elegant derivation; it is also a working tool used in weather forecasting, in industrial gas handling, in the design of insulating windows, and in the physiology of breathing. The final subtopic of this chapter gathers these applications together and shows kinetic theory at work in the world you live in.

APPLICATIONS OF THE KINETIC THEORY OF GASES

The kinetic theory of gases is not confined to physics examinations and laboratory demonstrations. It is a working tool, used every day by meteorologists explaining the weather, by doctors treating patients with respiratory problems, by engineers designing gas cylinders and window insulation, and by astronomers working out the chemistry of stellar atmospheres. Every phenomenon in which gas molecules do something visible or measurable can be traced back, eventually, to the handful of ideas developed in this chapter: molecules move rapidly and randomly, their speeds depend on temperature and mass, and the pressure they exert is the collective drumming of their collisions. This section takes six applications, each drawn from a different corner of life, and shows the kinetic theory at work in the world we live in.

1. Why atmospheric pressure falls with altitude

The atmosphere is held to the Earth by gravity. Each layer of air must support the weight of all the air above it, and since air at sea level has the greatest amount of atmosphere above it, sea-level pressure is the highest. As you climb a mountain, the amount of atmosphere above you decreases, and so does the pressure. At the top of Mount Kilimanjaro (5895m) the atmospheric pressure is about half its sea-level value; at the cruising altitude of a passenger jet (about 11000m) it is about one-quarter; at the altitude of the International Space Station there is essentially no atmosphere at all.

The molecular picture makes this quantitative. At any altitude, the local pressure equals the weight per unit area of the air column above. This weight is set by the total number of molecules in the column, which falls roughly exponentially with height because *faster-moving molecules can reach greater altitudes against gravity while slower ones cannot.* The result is that the number density of air molecules, N/V, falls with altitude in parallel with the pressure. A climber at the summit of Kilimanjaro draws about 54% as many air molecules into each breath as a climber at sea level. This is the number-density consequence of a pressure change, and it is why mountaineers at altitude struggle to breathe even though the chemical composition of the air (approximately 78% nitrogen and 21% oxygen) is identical at the summit and at sea level.

2. Gas cylinders and their safety

Domestic liquid petroleum gas (LPG) cylinders, widely used for cooking in Tanzanian homes, typically hold gas at pressures of several hundred kilopascals at room temperature. Industrial gas cylinders (oxygen, acetylene, nitrogen, argon) commonly hold gas at pressures of 10 to 20 megapascals, approximately 100 to 200 atmospheres. The walls of these cylinders must withstand the pressure for years without failure; the engineering rests on the physics developed in this chapter.

A sealed cylinder is a rigid container. Its volume does not change, and the amount of gas inside is fixed. From the ideal gas equation with V and n constant, the pressure is proportional to the absolute temperature: if the temperature rises, so does the

pressure, in direct proportion on the Kelvin scale. A cylinder left in direct tropical sunlight, as in the HOT example of an earlier section, can see its temperature rise from 20°C (293K) to around 70°C (343K), a 17% increase on the Kelvin scale. Its internal pressure rises by the same 17% fraction. A cylinder rated to hold gas safely at 12MPa in the morning can therefore see its internal pressure rise to over 14MPa by midday, pushing it close to or beyond its pressure rating.

This is the physical reason Tanzanian safety regulations require LPG and industrial gas cylinders to be stored in shaded, well-ventilated areas. The prohibition is not arbitrary bureaucracy; it is the direct application of the Kelvin-temperature scaling of gas pressure inside a fixed volume. Molecules inside a heated cylinder are moving faster and hitting the walls harder, and if the walls cannot bear the increased battering, they can rupture. A ruptured high-pressure cylinder releases its stored compressed gas explosively and has been responsible for industrial accidents across the world. The kinetic theory explains both why such accidents happen and how to prevent them.

3. Argon in double-glazed windows

Modern energy-efficient windows use two panes of glass separated by a thin sealed gap, typically a centimetre or two wide. The gap reduces heat loss through the window because still air is a poor thermal conductor. Filling the gap with argon instead of air reduces the heat loss further, and argon-filled windows are now standard in high-efficiency construction despite being more expensive than air-filled ones.

Two pieces of kinetic-theory physics explain why argon insulates better than air. Argon is a monatomic gas with only three degrees of freedom per atom, compared with five for the diatomic molecules nitrogen and oxygen that dominate air. Argon atoms therefore carry less thermal energy per collision than air molecules, which reduces the rate at which heat is transported across the gap. Additionally, argon atoms are heavier than the average air molecule (40g/mol versus about 29g/mol for air), so at the same temperature they move more slowly (v_{rms} proportional to $1/\sqrt{M}$), further slowing the transport of heat. Both effects act in the same direction, making argon a better thermal insulator than air by about 30%. The extra cost of argon is recovered over the lifetime of the window through reduced heating and cooling bills.

4. Diffusion of gases: Graham's law

When two gases are brought into contact, they spontaneously intermix until the mixture is uniform. This process, called diffusion, is driven entirely by molecular motion: molecules of each gas wander randomly through the other, and after enough time the concentrations become uniform. The rate at which a gas diffuses depends on how fast its molecules are moving, which in turn depends on their molar mass. Since $v_{rms} = \sqrt{(3RT/M)}$, lighter gases diffuse faster than heavier ones at the same temperature, and the ratio of diffusion rates is inversely proportional to the square root of the ratio of molar masses:

$$\frac{\text{rate}_1}{\text{rate}_2} = \sqrt{\frac{M_2}{M_1}}$$

This is Graham's law of diffusion, first stated empirically by the Scottish chemist Thomas Graham in 1848, long before the kinetic theory gave it a molecular explanation. It tells us, for example, that hydrogen (molar mass 2) diffuses through a porous membrane four times faster than oxygen (molar mass 32), because $\sqrt{(32/2)} = \sqrt{16} = 4$. Industrial applications include the enrichment of uranium for nuclear reactors, which relies on the very small difference in diffusion rates between the two uranium isotopes U-235 and U-238, separated through thousands of stages of gaseous diffusion. A small molecular-mass difference, amplified through many stages, produces the enriched product.

5. The speed of sound in gases

Sound is a pressure wave in a gas, propagated by successive collisions between molecules. A disturbance at one point compresses the local gas slightly, which pushes neighbouring molecules, which in turn push their neighbours, and so on, carrying the compression outward at a finite speed. It is plausible (and correct) that the speed of sound in a gas is close to the average speed of the molecules themselves, since the molecules are the carriers of the disturbance. A full derivation, accounting for adiabatic compression effects not yet treated in this book, gives:

$$v_{sound} = \sqrt{\frac{\gamma RT}{M}}$$

where γ is a number close to 1.4 for diatomic gases like air. Comparing this with the rms speed of the molecules, $v_{rms} = \sqrt{(3RT/M)}$, we see that the speed of sound is only slightly smaller than v_{rms}, differing by the factor $\sqrt{(\gamma/3)} \approx 0.68$. This agrees with experiment: in air at room temperature, the speed of sound is about 343m/s, while v_{rms} for air molecules is about 500m/s.

Two consequences follow.

- First, sound travels faster in lighter gases: the speed of sound in helium is nearly three times that in air, because helium molecules are lighter (M = 4 instead of 29) and therefore moving faster at the same temperature. This is why inhaling helium from a balloon makes your voice temporarily higher-pitched: the vocal cords vibrate at the same rate, but the sound waves in the helium-filled vocal tract travel faster, so the resonances shift to higher frequencies.

- Second, sound travels faster in warmer gases: the speed of sound in air rises by about 0.6m/s for each kelvin of temperature increase. This is why distant sounds carry further on cool evenings than on hot afternoons: the cooler air refracts sound downward, bending it along the ground rather than letting it escape upward.

6. Breathing and gas exchange

Breathing moves a mixture of gases (nitrogen, oxygen, carbon dioxide, water vapour, and trace gases such as argon) into and out of the lungs. Each of these gases, while mixed with the others, exerts its own partial pressure, contributing independently to the total pressure of the mixture. This is Dalton's law of partial pressures, met in the mixture example of the previous section. In the lungs, oxygen diffuses from the air into the blood, and carbon dioxide diffuses from the blood into the air, each following its own partial-pressure gradient. At sea level, the partial pressure of oxygen in inhaled air is about 21kPa; in the blood returning from the body, it is about 5kPa; the 16kPa difference is what drives oxygen into the blood.

At high altitude, total atmospheric pressure is lower, so the partial pressure of oxygen is also lower even though the percentage of oxygen in air has not changed. A climber at the summit of Kilimanjaro inhales air whose oxygen partial pressure is only about 10kPa. The partial-pressure gradient between lung air and returning blood is roughly halved, and the rate at which oxygen enters the blood is reduced correspondingly. This is the molecular reason for altitude sickness, and it is why climbers at extreme altitudes (above 7000m) typically use supplemental oxygen, which restores the partial-pressure gradient by raising the oxygen concentration in the inhaled gas.

The same physics operates in reverse for scuba divers. At 30m depth below the sea surface, the pressure on a diver is four atmospheres (one from the atmosphere plus three from the water column). Every gas the diver breathes has its partial pressure multiplied by four. Nitrogen, which is inert at sea-level partial pressures, becomes physiologically active at high partial pressure and can cause "nitrogen narcosis," a condition resembling drunkenness. Deep technical divers breathe gas mixtures (such as helium-oxygen) specifically chosen to reduce the partial pressures of problem gases, designing the breathing gas on the basis of Dalton's law.

With these six applications now in place, let us enjoy the ideas one last time through a pair of carefully chosen worked examples.

REAL Example 20

Explain for the following:

(a) A climber at the summit of Mount Kilimanjaro struggles to breathe even though the chemical composition of the air is the same as at sea level.
(b) Inhaling helium from a balloon temporarily raises the pitch of a person's voice.
(c) Hydrogen gas diffuses through a porous membrane four times faster than oxygen gas at the same temperature.

Solution

(a) At the summit of Kilimanjaro, the total atmospheric pressure is about half its sea-level value, so the partial pressure of oxygen in the inhaled air is also about half. The partial-pressure gradient that drives oxygen from the lungs into the blood is reduced, lowering the rate at which oxygen enters the bloodstream even though the oxygen fraction in the air (about 21%) has not changed.
(b) The speed of sound in a gas is proportional to $1/\sqrt{M}$, so it is higher in helium ($M = 4$g/mol) than in air ($M \approx 29$g/mol). When the vocal tract is filled with helium, sound waves travel faster through it, and the resonant frequencies of the tract shift upward. The vocal cords vibrate at their normal rate, but the filtering action of the tract emphasises higher frequencies, making the voice sound higher-pitched.
(c) Graham's law states that the rate of diffusion of a gas is proportional to $1/\sqrt{M}$. The ratio of diffusion rates for hydrogen ($M = 2$) to oxygen ($M = 32$) is $\sqrt{(32/2)} = \sqrt{16} = 4$. Hydrogen molecules are sixteen times lighter than oxygen molecules, so at the same temperature they move four times faster on average, and they diffuse four times faster through the same membrane.

Making Sense of the Answer: *three apparently different phenomena, one at altitude, one in a party balloon, and one in an industrial separation plant, all turn out to be applications of the same family of ideas: molecular speeds depend on temperature and mass, and physical processes that depend on molecular speed scale in predictable ways. This is the hallmark of kinetic theory: one set of principles, applied across a wide range of scales from the lung to the atmosphere to the uranium enrichment plant.*

Think Like a Physicist: *whenever an application question involves how fast a gas does something (diffuses, mixes, escapes, conducts sound), the first question to ask is how its molecular speed depends on the variables given. If the temperature is fixed, the answer almost always involves $1/\sqrt{M}$. If the gas is specified but the temperature varies, the answer involves $\sqrt{T}$. Recognising which scaling applies, and recognising it quickly, turns most application questions into a single ratio calculation.*

HOT Example 21

An industrial argon gas cylinder of internal volume 0.050m³ is filled with argon (molar mass 40g/mol) at 20°C to a pressure of 15MPa. The cylinder is rated to withstand pressures up to 18MPa. The cylinder is accidentally left standing in direct sunlight, where its temperature rises to 55°C.

(a) Calculate the new pressure inside the cylinder.
(b) State whether the cylinder is still within its safe pressure rating. Explain.
(c) Calculate the root mean square speed of the argon atoms at the higher temperature.

Solution

$T_1 = 20 + 273 = 293\text{K}; T_2 = 55 + 273 = 328\text{K}.$
(a) The cylinder is rigid (V constant) and sealed (n constant), so from $PV = nRT$ the pressure is proportional to temperature on the Kelvin scale:

$$P_2 = P_1 \times \frac{T_2}{T_1} = 15\text{MPa} \times \frac{328}{293} = 16.8\text{MPa}$$

The new pressure inside the cylinder is approximately 16.8MPa.
(b) The cylinder is rated to 18MPa. The new pressure of 16.8MPa is below the rating, so the cylinder is still within its safe operating range, though with a reduced safety margin of only about 1.2MPa. Continued heating of the cylinder could push the pressure past the rating and into the failure regime, so the cylinder should be moved to a shaded area promptly as a precaution.

(c) The root mean square speed of argon atoms at 328K uses $v_{rms} = \sqrt{\frac{3RT}{M}}$:

$$v_{rms} = \sqrt{\frac{3 \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (328\text{K})}{40 \times 10^{-3}\text{kg/mol}}} = 452\text{m/s}$$

The argon atoms have a root mean square speed of approximately 452m/s at 55°C.

Making Sense of the Answer: *the pressure rise of about 12% for a temperature rise of 35K is consistent with the direct Kelvin proportionality of the ideal gas equation at constant volume. The cylinder survives on this particular afternoon, but only with a small safety margin, and a cylinder rated closer to its initial operating pressure could easily be pushed past its limit by the same amount of solar heating. Cylinder safety ratings include temperature allowances for exactly this reason, and the recommended storage temperature is well below the conditions reachable in direct tropical sunlight.*

Think Like a Physicist: *sealed-cylinder problems at different temperatures are Kelvin-scale ratio problems. Convert both temperatures to kelvin, take the ratio, and apply it to the pressure. The molar mass, molecular speed, and other molecular properties do not enter the pressure calculation at all (they would only enter if the question asked about those properties); they are needed only when the question asks about speeds or kinetic energies separately. Identifying which quantities are needed for which part of the question saves time in a multi-part examination question.*

This chapter has taken the kinetic theory of gases from six assumptions to pressure, to temperature, to internal energy, and finally to six applications drawn from the world around you. Molecular motion is the common thread; the equations that follow from assuming many rapidly-moving particles bouncing elastically off each other turn out to describe the air in your lungs, the gas in your cooking cylinder, the insulation in your windows, the sound of your voice, and the weather above your head. The miscellaneous worked examples that follow bring all of these ideas together at once.

MISCELLANEOUS WORKED EXAMPLES ON THE KINETIC THEORY OF GASES

Example 22

(a) A gas in a sealed container exerts a steady pressure on the walls, even though the total momentum of all the molecules inside is zero. Explain how this is possible.
(b) A container of volume $5.0\times10^{-3}\text{m}^3$ holds 3.0×10^{24} molecules of a gas. Each molecule has a mass of $3.3\times10^{-26}\text{kg}$, and the mean square speed of the molecules is $1.6\times10^{5}\text{m}^2\text{s}^{-2}$. Calculate the pressure exerted by the gas on the walls of the container.

Solution

(a) Pressure arises from the change of momentum of individual molecules when they collide with the walls, not from any net momentum carried by the gas as a whole. A molecule approaching a wall reverses its component of momentum perpendicular to that wall, and the wall experiences an equal and opposite impulse. Because molecules are moving in all directions, the total vector momentum of the gas remains zero, but the rate of momentum change at each wall is non-zero and equal in magnitude on all walls. Pressure is the time-averaged force per unit area from this continuous rain of momentum reversals, so it does not require any net drift of the gas.

(b) Applying the kinetic equation $P = \frac{Nm}{3V}\bar{c}^2$:

$$P = \frac{(3.0 \times 10^{24}) \times (3.3 \times 10^{-26}\text{kg}) \times (1.6 \times 10^5\ \text{m}^2\text{s}^{-2})}{3 \times (5.0 \times 10^{-3}\ \text{m}^3)} = 1.06 \times 10^6\text{Pa}$$

The pressure on the walls of the container is approximately 1.06×10^6Pa.

Example 23

(a) Real gases obey the ideal gas equation very closely under some conditions but deviate from it under others. Explain the conditions under which a real gas departs most noticeably from ideal behaviour and identify which assumption of the kinetic theory fails first under those conditions.

(b) The molecules in a sample of gas are observed, and their speeds are recorded as follows:
3 molecules at 2.2m/s, 7 molecules at 6.1m/s, 5 molecules at 7.8m/s, 2 molecules at 4.0m/s, and 1 molecule at 2.5m/s. Calculate the root mean square speed of the sample.

Solution

(a) A real gas deviates most strongly from the ideal gas equation at high pressure and at low temperature.
- At high pressure, the volume of the molecules themselves is no longer negligible compared with the volume of the container, and the assumption of negligible molecular volume fails first.
- At low temperature, the weak attractive forces between molecules have time to act on the slowly-moving molecules, and the assumption of negligible intermolecular forces fails.

(b) The root mean square speed is obtained by taking the average of the squares of the speeds and then taking the square root. The mean of the squares is:

$$\bar{c}^2 = \frac{3 \times (2.2\text{m/s})^2 + 7 \times (6.1\text{m/s})^2 + 5 \times (7.8\text{m/s})^2 + 2 \times (4.0\text{m/s})^2 + 1 \times (2.5\text{m/s})^2}{3 + 7 + 5 + 2 + 1}$$

$$\bar{c}^2 = \frac{(14.52 + 260.47 + 304.20 + 32.00 + 6.25)\ \text{m}^2\text{s}^{-2}}{18} = 34.30\ \text{m}^2\text{s}^{-2}$$

$$v_{rms} = \sqrt{\bar{c}^2} = \sqrt{34.30\ \text{m}^2\text{s}^{-2}} = 5.86\text{m/s}$$

The root mean square speed of the molecules in the sample is approximately 5.86m/s.

Example 24

(a) Two sealed containers hold different gases at the same temperature. It is observed that although the average kinetic energy per molecule is the same in both containers, the average molecular speeds are different. Explain this observation.

(b) 2.0 moles of oxygen gas are contained at a temperature of 300K. Calculate the total translational kinetic energy of the gas sample, and comment on the magnitude of your answer by comparing it with the kinetic energy of a cricket ball of mass 0.16kg bowled at 36m/s.

Solution

(a) The average kinetic energy per molecule of an ideal gas depends only on the absolute temperature, according to the relation KE = (3/2)kT. Since the two gases are at the same temperature, they have the same average kinetic energy per molecule. However, the molecules of the two gases have different masses, so for the same kinetic energy $\frac{1}{2}m\bar{c}^2$ they must have different mean square speeds. The lighter gas has the higher rms speed and the heavier gas the lower rms speed, such that the product $m\bar{c}^2$ is the same for both.

(b) Total translational KE of the gas:

$$KE_{total} = \frac{3}{2}nRT = \frac{3}{2} \times (2.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K}) = 7483\text{J}$$

Kinetic energy of the cricket ball:

$$KE_{ball} = \frac{1}{2}mv^2 = \frac{1}{2} \times (0.16\text{kg}) \times (36\text{m/s})^2 = 104\text{J}$$

The total translational KE of 2.0mol of oxygen at 300K is approximately 7483J, which is about seventy-two times the kinetic energy of the cricket ball. This shows that a modest quantity of gas at room temperature carries substantial molecular kinetic energy, distributed across an astronomical number of molecules.

Example 25

(a) Explain what is meant by a degree of freedom of a gas molecule, and justify why a nitrogen molecule has five accessible degrees of freedom at room temperature while a helium atom has only three.

(b) 2.0 moles of helium gas and 2.0 moles of nitrogen gas are held in separate sealed rigid containers, both initially at 300K. Each sample is heated at constant volume and its temperature rises by 40K. Calculate the heat supplied to each sample, and compare the two values.

Solution

(a) A degree of freedom is an independent way in which a molecule can store kinetic energy. A helium atom, being essentially a point mass with no distinguishable axis of rotation, can store kinetic energy only in its translational motion along three perpendicular directions, giving three degrees of freedom. A nitrogen molecule, being a dumbbell shape, has the same three translational degrees of freedom but can also rotate about two axes perpendicular to its bond, adding two further degrees of freedom. Rotation about the bond axis itself stores negligible energy because the atoms lie essentially on the axis and the moment of inertia about it is very small, so it is not counted. Nitrogen therefore has f = 5 accessible degrees of freedom at room temperature.

(b) At constant volume the heat supplied equals the increase in internal energy of the gas:

$$Q = \Delta U = \frac{f}{2} \times nR\Delta T$$

For helium (f = 3):

$$Q_{He} = \frac{3}{2} \times (2.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (40\text{K}) = 998\text{J}$$

For nitrogen (f = 5):

$$Q_{N_2} = \frac{5}{2} \times (2.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (40\text{K}) = 1663\text{J}$$

Nitrogen requires approximately 1663J and helium approximately 998J to produce the same 40K temperature rise. Nitrogen needs about 1.67 times more heat because its molecules have two additional rotational degrees of freedom that must also be filled with energy as the gas warms; the ratio is exactly 5/3, the ratio of the degrees of freedom.

Example 26

(a) A student measures the temperature of a cold room as 10°C and the temperature of a warm room as 30°C. The student claims that the warm room is therefore three times as hot as the cold room. Explain why this claim is incorrect and state the correct ratio.

(b) Estimate the total number of air molecules in a classroom of dimensions 12m×8m×3.5m at a temperature of 25°C and a pressure of 1.0×10^5Pa. Use $k = 1.381 \times 10^{-23}\text{JK}^{-1}$.

Solution

(a) The Celsius scale places its zero at the freezing point of water, which is not the absolute zero of molecular motion. Ratios of temperatures measured on the Celsius scale are therefore not ratios of any physically meaningful quantity. The kinetic energy per molecule is proportional to the Kelvin temperature, not to the Celsius temperature. Converting: 10°C = 283K and 30°C = 303K. The correct ratio of hotness (measured by average molecular kinetic energy) is 303/283 ≈ 1.07, so the warm room is only about 7% hotter than the cold room, not three times as hot.

(b) Volume of the classroom:

$$V = 12\text{m} \times 8\text{m} \times 3.5\text{m} = 336\text{m}^3$$

Temperature in kelvin: $T = 25 + 273 = 298\text{K}$. Applying $PV = NkT$ to find N:

$$N = \frac{PV}{kT} = \frac{(1.0 \times 10^5\text{Pa}) \times (336\text{m}^3)}{(1.381 \times 10^{-23}\text{JK}^{-1}) \times (298\text{K})} = 8.2 \times 10^{27}$$

The classroom contains approximately 8.2×10^{27} air molecules.

Example 27

(a) Explain the molecular origin of the latent heat of fusion of a solid, using ideas from the kinetic theory of gases.

(b) A sealed cylinder contains a fixed mass of oxygen gas initially at a temperature of 20°C and a pressure of 2.0×10^5Pa. The gas is compressed to half its original volume and simultaneously heated to a new temperature of 100°C. Calculate the new pressure of the gas, the new root mean square speed of the oxygen molecules, and the new average translational kinetic energy per molecule. (Take the molar mass of oxygen to be 32g/mol, $k = 1.381 \times 10^{-23}\text{JK}^{-1}$.)

Solution

(a) In a solid, molecules are held in fixed positions by strong attractive intermolecular forces and vibrate about those fixed positions with a small kinetic energy. When heat is supplied at the melting point, that heat does not raise the kinetic energy of the molecules (the temperature does not change) but instead does work against the intermolecular forces, increasing the potential energy of the molecular arrangement. When enough heat has been supplied to weaken the forces that hold the solid together, the molecules become free to move past each other and the solid has melted into a liquid. The latent heat of fusion is the energy required per unit mass to perform this separation without any accompanying rise in temperature.

(b) Convert temperatures: $T_1 = 20 + 273 = 293\text{K}$; $T_2 = 100 + 273 = 373\text{K}$.
Applying the combined gas law $\frac{P_1V_1}{T_1} = \frac{P_2V_2}{T_2}$ with $V_2 = \frac{1}{2}V_1$:

$$P_2 = P_1 \times \frac{V_1}{V_2} \times \frac{T_2}{T_1} = (2.0 \times 10^5 \text{Pa}) \times 2 \times \frac{373\text{K}}{293\text{K}} = 5.09 \times 10^5 \text{Pa}$$

The new pressure of the gas is 5.09×10^5Pa.

New rms speed from $v_{rms} = \sqrt{\frac{3RT}{M}}$, at $T_2 = 373$K:

$$v_{rms} = \sqrt{\frac{3 \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (373\text{K})}{32 \times 10^{-3}\text{kg/mol}}} = 539\text{m/s}$$

The root mean square speed of the oxygen molecules at 373K is 539m/s.

Average translational KE per molecule at $T_2 = 373$K:

$$\text{KE} = \frac{3}{2}\text{kT} = \frac{3}{2} \times (1.381 \times 10^{-23}\text{JK}^{-1}) \times (373\text{K}) = 7.73 \times 10^{-21}\text{J}$$

The average translational kinetic energy per oxygen molecule at the new temperature is 7.73×10^{-21}J.

Example 28

(a) Two non-reacting ideal gases are mixed in a single container. Account for the fact that the total pressure exerted by the mixture on the walls of the container equals the sum of the pressures that each gas would exert alone in the same container at the same temperature.

(b) A sealed rigid container of volume 10.0 litres contains a mixture of 0.40 mol of helium gas and 0.60 mol of nitrogen gas at a temperature of 27°C. Calculate the total pressure of the mixture in the container and the total internal energy of the mixture. (Take helium to be monatomic and nitrogen to be diatomic.)

Solution

(a) In an ideal gas, intermolecular forces are negligible except during collisions, so the molecules of each species move independently of the molecules of the other species. Each species therefore collides with the walls of the container at a rate determined only by its own number density, molecular mass, and the temperature, unaffected by the presence of the other species. The pressure exerted by each species is the pressure it would exert if it alone occupied the container at the same temperature, which is called its partial pressure. The total force per unit area on the wall is the sum of the forces contributed by each species individually, so the total pressure equals the sum of the partial pressures. This statement is called Dalton's law of partial pressures.

(b) Convert: $V = 10.0\text{ litres} = 10.0 \times 10^{-3}\text{ m}^3 = 1.0\times10^{-2}\text{m}^3$; $T = 27 + 273 = 300$K.
Total moles: $n = n_{He} + n_{N_2} = 0.40 + 0.60 = 1.0$mol.
Total pressure from $PV = nRT$:

$$P_{tot} = \frac{nRT}{V} = \frac{(1.0\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K})}{1.0 \times 10^{-2}\text{ m}^3} = 2.49 \times 10^5 \text{Pa}$$

The total pressure of the gas mixture is 2.49×10^5Pa.

Total internal energy = sum of internal energies of each component.

Helium has $f = 3$ and nitrogen has $f = 5$:

$$U_{tot} = \frac{3}{2}n_{He}\,RT + \frac{5}{2}n_{N_2}\,RT$$

$$U_{tot} = \frac{RT}{2}(3n_{He} + 5n_{N_2})$$

$$U_{tot} = \frac{(8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (300\text{K})}{2} \times (3 \times 0.40\text{mol} + 5 \times 0.60\text{mol}) = 5236\text{J}$$

The total internal energy of the gas mixture is 5236J.

Example 29

(a) The kinetic theory of gases uses the root mean square speed (v_{rms}) of the molecules rather than their simple arithmetic mean speed. Explain the physical reason for this choice.

(b) A sample of gas contains 10 molecules with the following speeds: three molecules at 300m/s, four molecules at 500m/s, two molecules at 700m/s, and one molecule at 1000m/s. Calculate the simple arithmetic mean speed of the molecules, and separately the root mean square speed. Comment on why the two values differ.

Solution

(a) The pressure of a gas on the walls of its container depends on the rate at which molecules transfer momentum to the walls, and the momentum transfer per collision is proportional to the molecular speed. However, the kinetic energy transferred to or from each molecule is proportional to the square of its speed. The quantities that determine pressure and temperature are therefore controlled by the average of the squares of the molecular speeds, not by the average of the speeds themselves. The root mean square speed, being the square root of the mean square, is the speed that enters the pressure and temperature

equations correctly and has the right dimensions. The simple mean speed, while conceptually simpler, does not appear in the fundamental equations of the kinetic theory.

(b) Simple arithmetic mean:

$$v_{mean} = \frac{3 \times (300\text{m/s}) + 4 \times (500\text{m/s}) + 2 \times (700\text{m/s}) + 1 \times (1000\text{m/s})}{10} = 530\text{m/s}$$

Root mean square speed:

$$\bar{c}^2 = \frac{3 \times (300\text{m/s})^2 + 4 \times (500\text{m/s})^2 + 2 \times (700\text{m/s})^2 + 1 \times (1000\text{m/s})^2}{10} = 3.25 \times 10^5\ \text{m}^2\text{s}^{-2}$$

$$v_{rms} = \sqrt{\bar{c}^2} = \sqrt{3.25 \times 10^5\ \text{m}^2\text{s}^{-2}} = 570\text{m/s}$$

The simple arithmetic mean is 530m/s, while the root mean square speed is 570m/s. The rms is larger because squaring gives extra weight to the faster molecules, and the subsequent square root does not cancel this weighting.

Note the following fact: *Whenever the distribution of speeds has any spread, the rms speed exceeds the simple mean speed; the two are equal only if all molecules move at the same speed.*

Example 30

(a) Explain the molecular mechanism by which the pressure of a gas in a sealed rigid container increases when the temperature of the gas is raised.

(b) A rigid gas cylinder is rated to withstand internal pressures up to 1.5×10^6Pa. The cylinder is filled with nitrogen gas at a temperature of 20°C to a pressure of 1.0×10^6Pa. Determine the maximum temperature to which the gas can be heated before the pressure reaches the cylinder's rating, and calculate the percentage increase in the root mean square speed of the nitrogen molecules between the original temperature and this maximum.

Solution

(a) At a higher temperature, the molecules of the gas have a larger mean square speed and therefore move faster. Each molecule striking a wall delivers a larger momentum change per collision (because of its greater speed) and also strikes the walls more often per second (because it crosses the container more rapidly between collisions). Both effects multiply together: the rate of momentum transfer to the walls, and hence the pressure, rises in direct proportion to the absolute temperature at constant volume.

(b) At constant V and constant n, $PV = nRT$ gives $\frac{P}{T} = \text{constant}$.

$T_1 = 20 + 273 = 293\text{K}$:

$$T_{max} = T_1 \times \frac{P_{max}}{P_1} = (293\text{K}) \times \frac{1.5 \times 10^6\text{Pa}}{1.0 \times 10^6\text{Pa}} = 440\text{K} = 167°\text{C}$$

The maximum safe temperature for the gas is 167°C.

Since $v_{rms} \propto \sqrt{T}$:

$$\frac{v_{rms_{max}}}{v_{rms_1}} = \sqrt{\frac{T_{max}}{T_1}} = \sqrt{\frac{440\text{K}}{293\text{K}}} = 1.225$$

The rms speed increases by a factor of 1.225, a percentage increase of 22.5%.

Example 31

(a) Two sealed rigid containers hold different gases at the same temperature. The same quantity of heat is supplied to each container at constant volume. One gas is monatomic and the other is diatomic. Explain which gas undergoes the larger temperature rise and why.

(b) Two separate rigid containers, one holding 3.0 moles of helium gas and the other holding 3.0 moles of nitrogen gas, are initially at a temperature of 27°C. A quantity of heat equal to 5000J is supplied to each container at constant volume. Calculate the final temperature of each gas, and the ratio of the temperature rises.

Solution

(a) The heat supplied at constant volume goes entirely into increasing the internal energy of the gas. A diatomic molecule has five degrees of freedom and stores energy in translational motion and in two modes of rotation, while a monatomic molecule has only three degrees of freedom and stores energy in translation alone. For a given amount of heat, the diatomic gas must distribute that energy across five channels instead of three, so the energy per degree of freedom (and hence the temperature rise) is smaller in the diatomic gas. The monatomic gas therefore undergoes the larger temperature rise.

(b) At constant volume, $Q = \Delta U = \frac{f}{2} \times nR\Delta T$, so $\Delta T = \frac{2Q}{fnR}$.

For helium ($f = 3$):

$$\Delta T_{He} = \frac{2 \times (5000J)}{3 \times (3.0mol) \times (8.314Jmol^{-1}K^{-1})} = 133.7K$$

For nitrogen (f = 5):

$$\Delta T_{N_2} = \frac{2 \times (5000J)}{5 \times (3.0mol) \times (8.314Jmol^{-1}K^{-1})} = 80.2K$$

Final temperatures:

$$T_{f_{He}} = (27 + 273)K + 133.7K = 433.7K = 160.7℃$$
$$T_{f_{N_2}} = (27 + 273)K + 80.2K = 380.2K = 107.2℃$$

Ratio of temperature rises:

$$\frac{\Delta T_{He}}{\Delta T_{N_2}} = \frac{133.7K}{80.2K} = \frac{5}{3}$$

The helium reaches 160.7℃ and the nitrogen reaches 107.2℃.
The ratio of temperature rise of helium to that of nitrogen is 5: 3.

Example 32

(a) Two ideal gases with different molar masses are at the same temperature. When placed in thermal contact, no net transfer of heat occurs between them. Explain why this is so, even though the rms speeds of the molecules in the two gases are not equal.

(b) For one mole of an ideal gas at a temperature of 400K, calculate: (i) the average translational kinetic energy of a single molecule, (ii) the total translational kinetic energy of the gas sample, and (iii) the root mean square speed of nitrogen molecules at this temperature (molar mass of nitrogen = 28g/mol).

Solution

(a) Thermal equilibrium between two gases is a statement about the average kinetic energy per molecule, not about the average molecular speed. At the same temperature, molecules of both gases carry the same average translational kinetic energy, given by (3/2)kT. When molecules of the two gases interact at the boundary between them, collisions between a fast-moving light molecule and a slower-moving heavy molecule on average transfer no net energy in either direction, because the average kinetic energies are already equal. No heat flows between the two gases because no energy imbalance exists at the molecular level, even though the molecular speeds differ.

(b)(i) Average translational kinetic energy per molecule:

$$KE = \frac{3}{2} kT = \frac{3}{2} \times (1.381 \times 10^{-23} JK^{-1}) \times (400K) = 8.29 \times 10^{-21} J$$

The average translational kinetic energy per molecule is 8.29×10^{-21}J.

(ii) Total translational KE of one mole:

$$KE_{total} = \frac{3}{2} nRT = \frac{3}{2} \times (1mol) \times (8.314Jmol^{-1}K^{-1}) \times (400K) = 4988J$$

The total translational kinetic energy of one mole at 400K is 4988J.

(iii) Rms speed of nitrogen molecules:

$$v_{rms} = \sqrt{\frac{3RT}{M}} = \sqrt{\frac{3 \times (8.314Jmol^{-1}K^{-1}) \times (400K)}{28 \times 10^{-3} kg/mol}} = 597m/s$$

The root mean square speed of nitrogen molecules at 400K is 597m/s.

Example 33

(a) Explain why the rotational kinetic energy of gas molecules contributes to the internal energy of a diatomic gas but does not contribute to the pressure it exerts on the walls of its container.

(b) 2.0 moles of nitrogen gas (diatomic) are held at a temperature of 27°C. Calculate: (i) the total translational kinetic energy of the gas sample, (ii) the rotational kinetic energy per molecule, (iii) the total rotational kinetic energy of the sample, (iv) the total internal energy of the sample.

Solution

(a) Pressure arises from the change of momentum of molecules when they collide with the walls of the container, and only the translational motion of the centre of mass of a molecule contributes to this momentum change. A rotating diatomic molecule carries angular momentum about its own centre of mass, but rotation does not change the position of the centre of mass and therefore does not change the molecule's translational momentum. The rotational kinetic energy is internal to the molecule and does not influence the momentum transferred to a wall. Internal energy, however, is the total kinetic energy of all the molecules in all their modes of motion, so it counts translational and rotational kinetic energies together. Rotational KE adds to internal energy but not to pressure.

(b) $T = 27 + 273 = 300K$.
(i) Total translational KE = (3/2)nRT:

$$KE_{trans} = \frac{3}{2} \times (2.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 7483J$$

The total translational kinetic energy is 7483J.
(ii) For a diatomic molecule, rotational degrees of freedom $f_{rot} = 2$.
By equipartition each rotational degree of freedom stores (1/2)kT of energy per molecule, so the rotational KE per molecule is:

$$KE_{rot,1} = f_{rot} \times \frac{1}{2} kT = 2 \times \frac{1}{2} \times (1.381 \times 10^{-23}JK^{-1}) \times (300K) = 4.14 \times 10^{-21}J$$

The rotational kinetic energy per molecule is 4.14×10^{-21}J.
(iii) Total rotational KE of the sample = $(f_{rot}/2)$nRT:

$$KE_{rot,total} = \frac{2}{2} \times nRT = nRT = (2.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 4988J$$

The total rotational kinetic energy of the sample is 4988J.
(iv) Total internal energy = translational KE + rotational KE = (f/2)nRT with $f = 5$:

$$U = \frac{5}{2} nRT = \frac{5}{2} \times (2.0mol) \times (8.314Jmol^{-1}K^{-1}) \times (300K) = 12471J$$

The total internal energy of the sample is 12 471J, which is indeed the sum of the translational (7483J) and rotational (4988J) contributions.
(v) Ratio of rotational to translational KE:

$$\frac{KE_{rot,total}}{KE_{trans}} = \frac{4988J}{7483J} = \frac{2}{3}$$

The ratio of rotational to translational kinetic energy is 2:3.

Understand: *This fraction equals exactly the ratio of rotational to translational degrees of freedom (2 versus 3), as the equipartition theorem requires.*

Example 34

(a) The speed at which a sound wave travels through a gas is close to, but not exactly equal to, the root mean square speed of the molecules of the gas. Explain why the two speeds are closely related but not identical.
(b) A sample of an ideal gas undergoes a change in which its volume is doubled and its pressure is raised to four times its original value. If the root mean square speed of its molecules before the change was 250m/s, calculate the root mean square speed after the change.

Solution

(a) A sound wave is propagated through a gas by successive collisions between molecules: a local compression is passed from molecule to molecule across the gas, and the speed at which this handoff occurs cannot exceed the speed at which the molecules themselves are moving. The speed of sound is therefore approximately equal to the rms speed of the molecules. The two are not exactly equal because a sound wave involves compression, which is an adiabatic process, and the full analysis gives $v_{sound} = \sqrt{\frac{\gamma RT}{M}}$, whereas $v_{rms} = \sqrt{\frac{3RT}{M}}$. The two formulas differ by the factor $\sqrt{(\gamma/3)}$, which for a diatomic gas with $\gamma = 1.4$ gives a sound speed about 68% of the rms speed; for a monatomic gas with $\gamma = \frac{5}{3}$ the ratio is $\sqrt{\frac{5}{9}} = 0.75$. The two speeds therefore track each other closely but are distinct.
(b) From the ideal gas equation $PV = NkT$, if V is doubled and P is quadrupled:

$$\frac{T_2}{T_1} = \frac{P_2V_2}{P_1V_1} = \frac{4P \times 2V}{P \times V} = 8$$

The new temperature is therefore 8 times the original temperature.
Since $v_{rms} \propto \sqrt{T}$:

$$\frac{v_{rms_2}}{v_{rms_1}} = \sqrt{\frac{T_2}{T_1}} = \sqrt{8} = 2.83$$

$$v_{rms_2} = 2.83 \times (250m/s) = 707m/s$$

The new root mean square speed of the molecules is 707m/s.

Example 35

(a) An ideal gas, strictly defined, cannot be liquefied however much it is cooled or compressed. Explain why the liquefaction of real gases is not in conflict with the kinetic theory but is actually evidence about one of its assumptions.

(b) A sealed flask contains a fixed mass of neon gas initially at a temperature of 20°C and a pressure of 1.2×10^5Pa. The flask and its contents are cooled slowly to −50°C. Calculate the percentage change in the pressure of the gas and the percentage change in the root mean square speed of the neon atoms, and comment on whether the total internal energy of the sample has changed.

Solution

(a) Liquefaction requires molecules to stick together, and molecules can stick together only if there are attractive forces between them. The kinetic theory of an ideal gas assumes these intermolecular forces to be zero, so an ideal gas cannot liquefy by definition. Every real gas, however, liquefies when cooled sufficiently, because real molecules do attract each other weakly through van der Waals forces. The existence of liquefaction in real gases is therefore direct evidence that the assumption of zero intermolecular forces, while an excellent approximation at ordinary conditions, is not strictly true, and it fixes the regime (low temperature, high density) where that assumption begins to fail.

(b) $T_1 = 20 + 273 = 293\text{K}$; $T_2 = -50 + 273 = 223\text{K}$.
At constant V and constant n, $P \propto T$:

$$\frac{P_2}{P_1} = \frac{T_2}{T_1} = \frac{223}{293} = 0.761$$

Fractional change in pressure:

$$\frac{P_2 - P_1}{P_1} = 0.761 - 1 = -0.239$$

The pressure decreases by a percentage of 23.9%.
Fractional change in rms speed, using $v_{rms} \propto \sqrt{T}$:

$$\frac{v_{rms_2}}{v_{rms_1}} = \sqrt{\frac{T_2}{T_1}} = \sqrt{0.761} = 0.872$$

$$\frac{v_{rms_2} - v_{rms_1}}{v_{rms_1}} = 0.872 - 1 = -0.128$$

The rms speed decreases by a percentage of 12.8%.

The total internal energy of the sample has decreased. Neon is monatomic (f = 3), so the internal energy is U = (3/2)nRT, which is directly proportional to the Kelvin temperature. Since T has fallen from 293K to 223K, the internal energy has fallen by the same fraction of 0.239 (or 23.9%). The energy removed has been transferred as heat to the cooling medium surrounding the flask; it has not disappeared, but it is no longer in the gas.

Fourteen miscellaneous worked examples have put molecules, pressure, temperature, and degrees of freedom to work on problems that blend the chapter's ideas together. Having enjoyed the full combination of ideas, it is time to sharpen our thinking; the Digging Deeper Exercise is ready in the next page.

DIGGING DEEPER EXERCISE 4

EXERCISE 4A: BINDER QUESTIONS

Question 1

A student claims that the pressure of a gas in a container increases when the gas is compressed because *"the molecules touch each other more often."* Identify two errors in the student's reasoning and state what really causes the pressure increase.

Question 2

(a) Explain why the kinetic theory assumes that molecules in an ideal gas move in straight lines between collisions rather than in curved paths.

(b) Explain why the kinetic theory assumes that the time spent in a collision is negligible compared with the time between collisions.

Question 3

A cubic container and a cylindrical container hold equal amounts of the same gas at the same temperature and pressure. Compare the average rate at which molecules strike a square centimetre of wall in the cube with the average rate at which molecules strike a square centimetre of the curved wall of the cylinder, and explain your reasoning.

Question 4

(a) A student states that doubling the Celsius temperature of a gas must double the average kinetic energy per molecule. Explain why this claim is wrong.

(b) Show by algebra that no pair of Celsius temperatures exists for which the student's rule in (a) gives the correct ratio of 2 for the kinetic energies. Then describe the limit in which the student's rule becomes approximately correct, and explain briefly why the approximation improves in that limit.

Question 5

One mole of an ideal gas is held in container A, and two moles of the same ideal gas are held in container B of the same volume and at the same temperature.

(a) Identify which of the following quantities is the same in both containers: the pressure, the total internal energy, the average kinetic energy per molecule, and the root mean square speed of the molecules.
(b) Identify which of the above quantities is twice as large in container B as in container A.
(c) Using your answers to (a) and (b), explain briefly the distinction between quantities that are intensive and quantities that are extensive.

Question 6

A student computes the simple arithmetic mean of the speeds of a sample of ten molecules and calls the result the root mean square speed. Explain the student's error and identify the one circumstance in which the student's result would happen to equal the true root mean square speed.

Question 7

(a) Explain why a diatomic molecule such as nitrogen has two rotational degrees of freedom rather than three, even though rotation about three perpendicular axes is geometrically possible.

(b) Explain why the corresponding argument in (a) does not apply to a polyatomic molecule such as water vapour, which is treated as having three rotational degrees of freedom.

Question 8

A student argues that compressing an ideal gas (reducing its volume while keeping its temperature constant) must increase the gas's internal energy. Identify the error in the student's claim by appealing to the result, established in this chapter, that the internal energy of an ideal gas depends only on its absolute temperature.

Question 9

A sealed container holds a real gas. The container is compressed to one-tenth of its original volume at constant temperature.

(a) State whether the pressure of the real gas rises by exactly a factor of ten, by more than ten, or by less than ten, and explain which assumption of the kinetic theory is responsible for your answer.
(b) The compressed gas is then cooled to a very low temperature. Describe the new phenomenon that may occur for the real gas but has no analogue in the ideal gas model, and explain briefly why the ideal gas model cannot accommodate it.

Question 10

The absolute temperature of a sample of an ideal gas is increased from 300K to 1200K. Calculate the factor by which the root mean square speed of its molecules changes, and identify the scaling law you have used.

EXERCISE 4B: REAL QUESTIONS

Question 11

Kipute notices that a gas cylinder left in direct afternoon sun in Dar es Salaam feels noticeably warmer to the touch than one kept in the shade. The shop owner says this is the reason the spare cylinders are always stored in a shaded store-room.

(a) Explain, using the molecular picture, why the pressure inside the warm cylinder has increased even though no gas has been added.

(b) Explain why Tanzanian safety regulations require cylinders to be stored in shaded, well-ventilated places rather than in closed rooms.

Question 12

At Horombo Huts on Mount Kilimanjaro (altitude 3720m), a mountain climber feels mildly breathless. By the time he reaches Uhuru Peak (altitude 5895m) he is gasping between every few words. His guide assures him that the air composition has not changed: it is still about 78% nitrogen and 21% oxygen. Explain to the mountain climber what physical quantity has in fact changed with altitude, and how this accounts for the increasing difficulty of breathing.

Question 13

A mother cooks beans in a sealed pressure cooker. The cooker's safety valve is set to release steam if the internal pressure exceeds twice atmospheric pressure.

(a) Explain, using the kinetic theory, why the pressure inside the sealed cooker rises as the contents are heated from below.

(b) Explain why the pressure reaches a stable value rather than rising indefinitely, even if the flame remains lit.

Question 14

A diver descends to a depth of 20 metres below the surface of the sea off Zanzibar, where the total ambient pressure is about three atmospheres. The diver is breathing ordinary compressed air.

(a) Explain why the partial pressure of nitrogen in the diver's lungs at this depth is approximately three times its value at sea level.

(b) Nitrogen is chemically inert at sea level but is known to induce symptoms resembling drunkenness ("nitrogen narcosis") at high partial pressures. Explain, using molecular collision rates, why a chemically inert gas can become physiologically active once its partial pressure is raised sufficiently.

Question 15

Ally claims that a helium balloon rises because *"helium atoms are lighter than air atoms."* **Bakari** objects that Ally is wrong because air is a mixture of different gases rather than a single kind of atom. Resolve the disagreement using the kinetic theory, and explain the sense in which Ally's statement captures the correct physics despite the careless wording.

Question 16

An architect designing an upmarket office in Arusha specifies double-glazed windows filled with argon gas rather than air, despite the higher cost of the argon-filled version.

(a) Explain, using kinetic theory, two distinct molecular properties of argon that make it a better insulator than air.

(b) Explain whether replacing argon with krypton (molar mass 84 g/mol, monatomic) would improve the insulation further, identifying which of the two properties changes in the favourable direction.

Question 17

Kipute pumps up her bicycle tyre on a hot Dodoma afternoon. She notices that the metal barrel of the pump becomes warm during vigorous pumping, even though she has not heated it with any external flame. Explain what happens to the air trapped in the pump during each downstroke of the handle, and account for the warmth of the pump body in molecular terms.

Question 18

Air in a shaded well in a rural village is at 24°C, while air on a Dar es Salaam street at midday is at 36°C. A student claims that *"the air molecules on the street move much faster than those in the well."* Evaluate this claim by computing the ratio of root mean square speeds in the two locations, and comment on whether the adverb "much" is appropriate.

EXERCISE 4C: HOT QUESTIONS

Useful constants: *Boltzmann constant $k = 1.381\times10^{-23}\ JK^{-1}$; Avogadro constant $N_A = 6.022\times10^{23}\ mol^{-1}$; universal gas constant $R = 8.314\ Jmol^{-1}K^{-1}$; standard atmospheric pressure $= 1.013\times10^{5}\ Pa$.*

Question 19

A steel cylinder of internal volume $0.040m^3$ contains nitrogen gas at 17°C and a pressure of 8.0×10^6Pa. A fault develops and the gas temperature rises to 127°C. Calculate the quantity of heat supplied to the gas during this temperature rise.

Question 20

Starting from the kinetic equation $PV = \frac{1}{3}Nm\bar{c}^2$ and the equation of state $PV = NkT$, derive the relation between the root mean square speed of the molecules of an ideal gas and its Kelvin temperature. Hence prove that for any two ideal gases A and B at the same temperature, the ratio of their rms speeds equals the square root of the inverse ratio of their molar masses, and explain briefly the physical reason why this ratio does not depend on temperature.

Question 21

A rigid cylinder holds 4.0 moles of neon gas at 27°C. It absorbs 2000J of heat at constant volume. A second rigid cylinder, identical in shape and size, holds 4.0 moles of oxygen gas also at 27°C and also absorbs 2000J at constant volume. Which cylinder ends up at the higher pressure, and by what factor?

Question 22

A pure gas is observed to diffuse through a porous membrane at a rate equal to one divided by 1.41 times the rate of hydrogen at the same temperature. Identify the gas and state briefly why your identification is unique among common gases at A-Level. (The molar mass of hydrogen is 2.0g/mol.)

Question 23

A rigid sealed container of volume $2.0\times10^{-3}m^3$ holds helium gas at 127°C and 5.0×10^5Pa. The container is placed in a room at 27°C and the helium inside cools according to Newton's law with cooling constant $\lambda = 2.0\times10^{-3}s^{-1}$. Find the pressure inside the container 10 minutes after it is placed in the room.

Question 24

A sealed vessel of volume 25 litres contains 1.5 moles of helium, 2.0 moles of nitrogen, and 0.5 moles of water vapour, all at 127°C. A dessicant is then introduced that completely absorbs the water vapour, removing it from the gas phase. Calculate the percentage change in the total pressure of the mixture and the percentage change in the total internal energy of the mixture, and explain why the two percentages are not equal.

Question 25

An ideal monatomic gas expands from volume V_0 to volume $2V_0$ along a process in which the pressure and volume are related by $P = \alpha V$, where α is a positive constant. The initial pressure is P_0 and the initial absolute temperature is T_0. Find the factor by which the root mean square speed of the molecules has changed by the end of the expansion, and find the final internal energy of the gas as a multiple of its initial internal energy.

Chapter 5

THERMODYNAMICS

INTRODUCTION

The aluminium kettle had been sitting on Mr. Akilikubwa's desk since before the first bell, plugged into the wall with the quiet confidence of a demonstration that has been planned for days. It was a small kettle; the kind a college student might keep in a hostel room for boiling water for instant noodles at eleven o'clock at night. But this one had a modification. Balanced on top of the lid, perfectly centred, was a 200-gram iron weight, the squat flat kind that sits on a kitchen scale when a *mama ntilie* is measuring out rice. It was not decorating the kettle. It was sitting there with purpose, the way things on Mr. Akilikubwa's demonstration bench always had purpose, even when that purpose was not yet visible.

Kipanga noticed the kettle the moment he walked in. He noticed it the way a man notices a parked car with its hazard lights on, not as a crisis exactly, but as something that will probably become relevant to him personally before the day is over.

He took his seat. He kept one eye on the kettle.

Mr. Akilikubwa opened the register, clicked the kettle on, and began calling names without looking up. The kettle was apparently not part of the lesson yet. It was simply a kettle on a desk, minding its own business, heating up.

Kipanga: *"Sir, the kettle is on."*

Mr. Akilikubwa: *"Yes, Kipanga."*

Kipanga: *"Should the kettle be on?"*

Mr. Akilikubwa: *"That is an interesting question and the answer is yes."*

Three minutes passed. The kettle began to hiss, a sound that started low and private and slowly became impossible to ignore. Kipute, who had been taking notes from the previous lesson, put her pen down. The class, by some collective instinct, turned to face the front.

Then the lid moved.

It did not fly off. It did not explode. It simply lifted, just slightly, perhaps two centimetres, tilted forward at a polite angle, and hovered there, trembling, as a thick tongue of steam curled out from underneath it. The 200-gram iron weight, which had been sitting perfectly flat, was now visibly higher than it had been when the lesson started. The steam was holding it up. Quietly, steadily, and without any drama, the steam inside that small aluminium kettle was pushing a 200-gram piece of iron upward against gravity.

The class stared.

Mr. Akilikubwa looked up from the register. He regarded the floating lid with the mild satisfaction of a man who has seen this before and is still, after all these years, quietly pleased that it works.

Mr. Akilikubwa: *"Kipanga."*

Kipanga: *"Sir."*

Mr. Akilikubwa: *"Please come and press the lid back down."*

Kipanga blinked. This seemed reasonable. The lid was misbehaving. He was being asked to correct it. He was a helpful person. He stood up, walked to the desk, extended his right index finger, and pressed down on the lid with the casual confidence of someone who has never thought carefully about what is underneath it.

For a long and instructive moment, Kipanga's finger and the steam inside the kettle conducted a quiet negotiation.

The steam won!

The lid pushed back. Not violently, not dangerously, but with a steady, warm, entirely undeniable upward force that Kipanga had not budgeted for. He pressed harder. The lid pushed back harder. He was, at this moment, a 200-gram iron weight with a university scholarship. He was losing to water vapour.

He stepped back. His finger was pink and slightly damp.

Kipanga: *"It pushed me."*

Mr. Akilikubwa: *"Yes."*

Kipanga: *"The steam pushed my finger. Upward."*

Mr. Akilikubwa: *"Yes."*

Kipanga: *"That is my hand."*

Mr. Akilikubwa: *"Also yes. Please sit down before you have any further conversations with the kettle."*

Kipanga sat down. He looked at his finger. He looked at the kettle. The lid was still hovering. The iron weight was still in the air. He had the expression of a man who has just been bested in arm-wrestling by a cup of boiling water and is not quite sure which part of his life philosophy this should prompt him to revise.

Kipute had not moved from her seat. She had not volunteered, had not rushed forward, had not touched anything. She had simply watched the entire exchange with the calm attention of someone assembling evidence.

Kipute: *"Kipanga, where did the energy come from? To push your finger. And to keep that weight in the air."*

Kipanga: *"From the steam."*

Kipute: *"And where did the steam get it?"*

Kipanga: *"From the hot water."*

Kipute: *"And where did the hot water get it?"*

Kipanga looked at the power cord running from the kettle to the wall socket. Then he looked at his finger again. Then he looked at the iron weight, still hovering above the lid.

Kipanga: *"From the electricity."*

Kipute: *"So electricity went in through that cable, and a 200-gram weight went up into the air. Something converted one into the other. What was the something?"*

Kipanga opened his mouth. He closed it. He was not used to being asked questions whose answers he did not immediately have. He was especially not used to Kipute being the one asking them.

At that moment Mr. Akilikubwa switched the kettle off. The hissing faded. The lid sank back with a soft aluminium clunk. The iron weight settled. The steam dissipated. The classroom returned to normal, except for the question still hanging in the air, which is exactly where Mr. Akilikubwa intended it to remain.

Mr. Akilikubwa: *"Kipute has asked the question that this entire chapter exists to answer. Electricity went in. A weight went up. Something in the middle did the converting. That something was a thermodynamic process. And the framework that lets us track every joule, every step of the way, from the wall socket to the iron weight and everywhere in between, is called the* ***first law of thermodynamics****."*

He pulled the power cord from the wall and set the kettle aside.

Mr. Akilikubwa: *"Chapter 4 told you what temperature is at the molecular level. It told you that the internal energy of a gas is the total kinetic energy of its molecules, and that the temperature of the gas is simply a measure of how fast those molecules are moving. Beautiful. Correct. But Chapter 4 had one quiet limitation that I never pointed out, because you were not ready for it. In Chapter 4, the gas sat inside its container and did its own molecular business. The walls never moved. The gas never pushed anything. No piston slid. No lid lifted. No finger was embarrassed in front of thirty witnesses."*

Kipanga stared at the kettle with fresh resentment.

Mr. Akilikubwa: *"This chapter takes the walls off the nail. Here, the gas is allowed to push on pistons, on lids, on the atmosphere, on whatever is in its way. And when it pushes something through a distance, it does work. Mechanical work. The same work you calculated in Form Four when you pushed a box along a floor. The same joules. The same equation. The difference is that here, the gas is doing the pushing, and the energy to do that pushing came from somewhere, from heat, or from compression, or from both at once. The first law is simply the rule that the accounts must balance. Every joule that goes into a gas must either raise its internal energy or come out again as work. Not some of the joules. Not most of the joules. Every single one."*

The class was quiet in the way that classes are quiet when something has genuinely landed.

Kipute wrote something in her notebook. Kipanga was still looking at his finger, but his expression had changed. It was no longer the expression of a man who had been beaten by a kettle. It was the expression of a man who was beginning to understand that the kettle had been trying to teach him something.

Chapters 2, 3, and 4 have been building toward this point from three different directions. Chapter 2 established that heat is energy in transit between bodies at different temperatures, gave us the Zeroth Law to make the word temperature logically meaningful, and developed the practical tools of specific heat capacity and latent heat that let us calculate how many joules it takes to boil a kettle, melt a block of ice, or cool a hospital ward. Chapter 3 traced the three routes heat can travel, through conduction, convection, and radiation, and showed that the rate of flow matters as much as the total amount. Chapter 4 then reached inside the gas itself, connected temperature to molecular kinetic energy, and delivered the single most important result so far: the internal energy of an ideal gas is $U = (f/2)nRT$, a number written entirely in terms of how many molecules are present and how fast they are moving.

In all three of those chapters, the gas was a passive participant. It sat inside its container, accumulated energy, increased its temperature, and was grateful. Nothing in those chapters required the gas to push on anything.

This chapter removes that restriction.

When the walls of a container are free to move, the gas can do mechanical work on its surroundings. The steam in Mr. Akilikubwa's kettle is pushing the lid upward. A gas in a car engine is pushing a piston downward. The hot gas in a turbine is pushing a blade sideways at hundreds of metres per second. In all of these, the same principle applies: the gas exchanges energy with its surroundings by two channels simultaneously. The first channel is heat, which is energy flowing across a temperature difference, as in Chapters 2 and 3. The second channel is work, which is energy transferred by a force acting through a distance, exactly as in mechanics. The first law of thermodynamics is the single equation that keeps the accounts for both channels at once: every joule entering the gas either increases its internal energy or leaves as work done on the surroundings. There is no third option, no hidden drawer, no energy that simply disappears. The books balance.

This chapter unfolds in three parts.

Welcome to Chapter 5. Kipanga's finger has already had its first lesson. The rest of the chapter will explain what it felt.

THE LANGUAGE OF THERMODYNAMICS

The kettle is back on the bench and the weight is sitting still, but Kipute's question has not gone anywhere. Electricity went in, a weight went up, and something in the middle did the converting. Before we can write an equation for what that something was, we need a small set of words. Thermodynamics is a subject where the vocabulary carries its own weight. Using "heat" when one means "temperature," or "system" when one means "surroundings," will send a student sideways for pages at a time. This section is the dictionary. There are no equations here and no worked examples. What there is, is a short list of terms that every subsequent section will assume has been understood.

Some of the vocabulary is already in place. Chapter 2 distinguished temperature from heat, defined thermal equilibrium, and gave us the **Zeroth Law** that makes the word temperature meaningful in the first place. Chapter 4 introduced the **internal energy U** of a gas and tied it to the temperature and to the molecular degrees of freedom. Those definitions are assumed from this point forward, and the reader is invited to turn back to the earlier chapters if any of that language has gone cold. What is new, and what this section adds, is the vocabulary we need to talk about a gas that is exchanging energy with the world outside it, which is to say, a gas that is doing something to Kipanga's finger.

System, surroundings, and boundary

When a physicist wants to think about some piece of the world, she starts by drawing an imaginary line around it. Everything inside the line is called the **system**. Everything outside the line is called the **surroundings.** The line itself is called the **boundary**.

Consider the kettle. If we are interested in what the steam did to the lid, the natural choice is to call the steam (together with the water underneath it, if we are being thorough) the system. The boundary is then the inner surface of the kettle, together with the underside of the lid. The surroundings are everything else in the story, namely the kettle walls themselves, the 200-gram weight, Kipanga's finger, the hot plate below, and the rest of the classroom. This choice is a decision, not a discovery. We could equally well call the kettle-and-its-contents the system, in which case the boundary becomes the outer surface of the kettle and the hot plate is surroundings. Neither choice is wrong. They are two different questions about the same physical object, and which one we prefer depends on what we are trying to calculate.

Throughout this chapter, the system will almost always be some quantity of gas inside some container, and the surroundings will be everything outside that container. The boundary will sometimes be rigid, like the walls of a sealed tank, and sometimes movable, like the upper face of a piston or the bottom of a hovering kettle lid. But the principle is general. The first thing one does in any thermodynamics problem is decide what the system is, and once that choice is made, one stays with it. Changing one's mind halfway through a question and calling the piston the system when the gas was the system a moment ago is how sign errors are born.

Open, closed, and isolated systems

Three things can, in principle, cross the boundary of a system.

1) Matter can cross, in the form of molecules moving from one side to the other.
2) Energy can cross as heat, which is the mode of transfer we met in Chapter 2.
3) Energy can also cross as work, which is the mode that appears whenever a force is applied through a distance.

Which of these three the boundary lets through determines what kind of system we are dealing with.

An **open system** *allows all three to cross.* A pot of *ugali* on a charcoal *jiko* with its lid off is open. Steam escapes into the air (matter leaves), the flame delivers heat through the base of the pot (heat enters), and a *mwiko* stirring the ugali does work on the mixture (work enters). Mr. Akilikubwa's kettle during the demonstration was also open, briefly, whenever the lid lifted and the steam escaped around its edges.

A **closed system** *allows heat and work to cross, but not matter.* The quantity of stuff inside is fixed. A sealed saucepan with a snug lid is closed. No steam escapes. Heat still flows in from the flame, and pressing on the lid still does work on the air trapped inside. The kettle, when its lid was resting flat and no steam was getting past it, was behaving as a closed system.

An **isolated system** *allows nothing at all to cross.* No matter, no heat, no work. An isolated system is hermetically sealed, perfectly rigid, and perfectly insulated from everything around it. The iconic example is the thermos flask of hot tea that sits on a desk and refuses to cool for hours. Its double-walled vacuum blocks conduction and convection, its silvered inner surface blocks radiation, and its screw-top blocks mass flow. No real flask is perfectly isolated, and the tea will eventually cool if given enough time, but for the duration of a physics problem the flask is close enough to isolated that we treat it as isolated.

Throughout this chapter we work almost entirely with closed systems. The first law of thermodynamics, which is the business of the next section, is a statement about closed systems. A gas sealed in a rigid tank is closed. A gas sealed beneath a movable piston is also closed, because the piston moves but the molecules stay inside. The kettle with its lid hovering was at the edge of closed and open, which is why the demonstration was interesting but also why a quantitative analysis would need us to be careful about the steam leaking out. For the next several sections we will work with containers that do not leak.

Thermodynamic state and state variables

Once we have decided what the system is, we need to describe its condition. The physical condition of a gas can be specified completely by a small number of measurable quantities: its **pressure P**, its **volume V**, its **temperature T**, and the **amount of substance n**, which we measure in moles. These four quantities are called **state variables.** They have this name because *a complete specification of all four amounts to a complete description of the gas.* Two samples of gas with the same P, V, T, and n are thermodynamically identical, whether one is in Bukoba and the other in Mbeya. Nothing else need be said about them. Their future behaviour, if handed over to another physicist, will be indistinguishable.

The set of values (P, V, T, n) at a particular moment is called the **thermodynamic state** of the system. The state is a snapshot, a list of numbers true at one instant. At the moment Mr. Akilikubwa switched on the hot plate, the air inside the kettle was in one state. By the time the lid lifted, it was in a different state: the temperature had risen, the pressure had built up, and some of the water had joined the gas as steam, so the amount of substance present in the gaseous part of the system had increased. Every step of the demonstration was a change from one state to another.

The Zeroth Law of Chapter 2 is quietly doing its work every time we write the letter T in that list. Without the Zeroth Law, the temperature of the gas would not be a single well-defined number for the whole body; it would be some tangle of pair-by-pair relationships, and the state description would fall apart before we got started. It does not fall apart, because the Zeroth Law holds.

For an ideal gas, the four state variables are not independent. They are tied together by the equation of state $PV = nRT$ that we derived in Chapter 4. Knowing any three of them determines the fourth. Specify P, V, and n, and T is forced. Specify P, T, and n, and V follows. The state space of an ideal gas is therefore really three-dimensional, even though we talk about four variables. In practice this is convenient, because most problems hold one of the four variables fixed while the others change, and the language of "constant pressure" or "constant temperature" is more natural than insisting on three degrees of freedom and solving the constraint each time.

Equilibrium and quasi-static processes

The state variables P, V, T have meaningful values only when the gas has had time to settle down, so that its molecules have redistributed themselves evenly, any temperature gradients have died out, and the pressure has equalised from wall to wall. The notion of thermal equilibrium between two separate bodies was defined in Chapter 2. Here we extend it to a single body in equilibrium with itself. A gas is in **thermodynamic equilibrium** when its state variables are well-defined, uniform throughout the body of gas, and unchanging with time in the absence of any external disturbance. If we have left a tank of gas alone for long enough, its pressure is the same at the top as at the bottom, its temperature is the same in the middle as at the walls, and if no one switches on a flame or pulls out a piston, all the state variables will still be the same an hour from now. That is equilibrium.

Thermodynamics is a subject about equilibrium states joined by processes. The gas starts in some equilibrium state A and finishes in some other equilibrium state B. The transition is, strictly speaking, complicated. During the rapid lifting of Mr. Akilikubwa's kettle lid, the pressure was not uniform across the gas inside the kettle; pressure waves were bouncing from wall to wall, and the temperature near the hot plate was different from the temperature at the top. For calculations, this is inconvenient. A gas whose state is not well defined cannot be assigned a point on a PV diagram, and a process whose intermediate stages cannot be drawn cannot have its work easily computed.

The standard way around this is to assume the process happens slowly enough that the gas passes through a continuous sequence of near-equilibrium states. A process that satisfies this condition is called **quasi-static**. Quasi-static is an idealisation, not a description of any real process, but it is an idealisation good enough to get the right answer to within a few percent in most A-Level problems, which is the accuracy we are aiming for. When a question describes the gas as "slowly compressed"

or "gradually heated," it is instructing us to treat the process as quasi-static. When a question says the tyre "suddenly bursts" or the gas is "released rapidly," it is flagging precisely the opposite, namely a non-quasi-static process that we handle differently. We will meet that different handling later in the chapter, when we reach the adiabatic process.

Thermodynamic process and the path that joins the states

A **thermodynamic process** *is any change of state*. The gas starts in state A and ends in state B, and the process is the sequence of intermediate states that takes it from one to the other. The demonstration with the kettle was a thermodynamic process, as was the cooling of the tea in the thermos flask, and as is every compression of the air in a bicycle pump, every expansion of the gases in a car engine, and every breath the reader has ever taken.

Here is the central idea that makes this subject interesting, and it is the idea that will be formalised in the section on work done by a gas. The same initial state A and the same final state B can be connected by many different processes, by many different paths through state space. Different paths do different things along the way. One path might heat the gas first and compress it afterwards. Another might compress it first and heat it afterwards. A third might do both at once. The final state is the same either way, but how much heat flowed in along the way, how much work was done, and in which direction each was pointing, all of these depend on the path. In thermodynamic language, **heat and work** are *path-dependent* quantities. The **state** itself is *not*. A snapshot does not care how you arrived at it. The bookkeeping does.

This is the deepest reason why the four processes we are about to meet give four different answers for the work done, even when they start and end in identical states. It is also the reason we draw PV diagrams, which are the tool for making path-dependence visible. For now, hold on to the shape of the idea. State is where you are; process is how you got there; and the "how" affects the balance sheet even when the "where" is the same.

The four workhorses

Thermodynamic processes come in many varieties, but four of them are worth particular study. They are simple, they are common, and between them they cover almost all the physics an A-level student needs. What distinguishes the four is which of the state variables is held constant during the process. The names, in the tradition of the subject, are Greek.

The four processes are collected in the preview table below. The heat Q and work W columns are left empty for now, because filling them in is the business of the sections that follow. Treat the table as a roadmap. Put a bookmark at this page. After the isochoric section is read, come back and fill in the isochoric row; after the isobaric section is read, fill in the isobaric row; and so on. By the end of the chapter the table will be complete, and the complete version will become the single most important reference in this chapter.

Process	What is constant	PV diagram shape	Heat Q	Work W
Isochoric	Volume V	Vertical straight line	to come	zero (no volume change)
Isobaric	Pressure P	Horizontal straight line	to come	to come
Isothermal	Temperature T	Gentle hyperbola, $PV = \text{constant}$	to come	to come
Adiabatic	Heat exchange $(Q = 0)$	Steep hyperbola, $PV^{\gamma} = \text{constant}$	zero (no heat flows)	to come

The Greek roots explain the names. **Isochoric** combines **iso** (*equal*) with **choros** (*space, volume*), so the gas is held at constant volume. **Isobaric** combines **iso** with **baros** (*pressure*), so the gas is held at constant pressure. **Isothermal** combines **iso** with **thermos** (*heat, temperature*), so the gas is held at constant temperature. And **adiabatic** comes from **adiabatos**, meaning *"not passing through".* It is a name about the heat, which in an adiabatic process does not pass through the walls of the container. Physicists have been borrowing words from Greek for thousands of years, and the habit is now too strong to change.

Each of these four processes is met in turn as the chapter unfolds. For each we draw the PV diagram, derive the work done, find how much heat flows, and compute the change in internal energy. By the time the synthesis section is reached, the reference table above will be complete in every cell. Before any of that can happen, however, we need the law at the heart of this subject, the law that answers Kipute's question about what converted the electricity into a hovering iron weight. That is the business of the next section.

THE FIRST LAW OF THERMODYNAMICS

At the end of the introduction, Kipute asked a question. Electricity had flowed in through a cable, a 200-gram iron weight had gone up two centimetres into the air, and something in the middle had done the converting. The previous section gave the something a name, a **thermodynamic process**, and put six words on the desk to talk about it: system, surroundings, boundary, state, equilibrium, path. Naming a thing is a good first step. It is not the same as understanding it. The rule that tracks every joule of Kipute's electricity, from the hot plate through the water and the steam and out again into the gravitational potential energy of a kitchen-scale weight, is called the **first law of thermodynamics**, and it is the single most important sentence in this chapter.

Return to the kettle for a moment. The hot plate is on, the lid is resting flat, and the steam inside is quietly building pressure. Three things are happening at once, and the first law is nothing more than the observation that these three things always balance. Heat is flowing into the kettle from the hot plate, because the hot plate is hotter than the water. The water and the steam inside the kettle are getting hotter, which means their internal energy is rising, because internal energy is just the total kinetic energy of the molecules, and hotter molecules move faster. And if the lid is free to move, some of that incoming energy is going back out again as mechanical work done on the lid, and on whatever is resting on the lid. Three things. Three arrows. One kettle. The books balance.

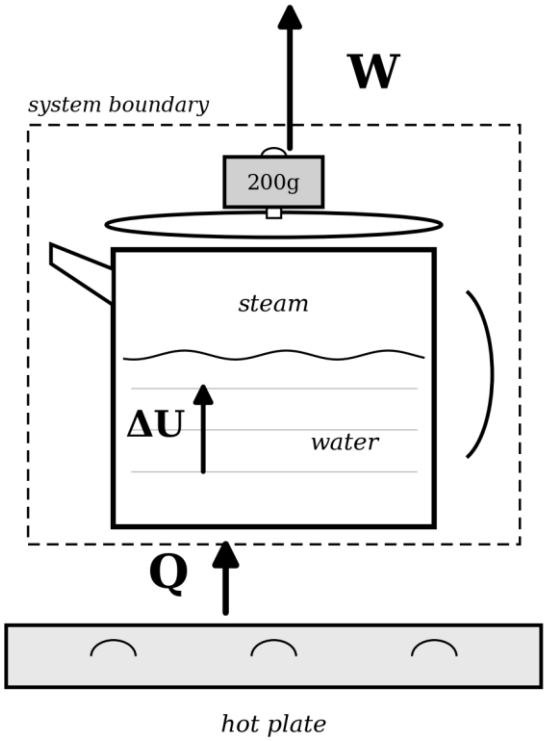


Figure: ***The three energy flows at work in Mr. Akilikubwa's kettle.*** *Heat Q enters from the hot plate below. Work W leaves through the rising lid (which is lifting the 200g iron weight against gravity). The change in internal energy ΔU, shown as a vertical arrow inside the body of water, tracks whatever is neither entering nor leaving. The first law says these three arrows balance to the last joule.*

The formal statement is short, and it is worth reading twice. *For any closed system undergoing a thermodynamic process, the heat Q supplied to the system is equal to the increase ΔU in its internal energy plus the work W done by the gas on its surroundings:*

$$Q = \Delta U + W$$

Every joule of heat that goes into the gas either stays inside the gas as added internal energy, or leaves again as mechanical work done on the surroundings, or splits itself between the two in some ratio. There is no third option. There is no hidden drawer in the wall of the kettle where a joule can disappear and reappear later. If a joule went in as heat, it is still there in one of the two forms, or it has already left again as work. That is the whole of the law, and that is why, at the instant Mr. Akilikubwa switched on the hot plate, the outcome of the entire demonstration was already decided. The electricity was going to go in. Some of it was going to stay inside the kettle as hotter water and hotter steam. The rest was going to come back out again as the gravitational potential energy of a rising iron weight. The exact split depended on the details, but the total was going to balance. Kipute's question, ten minutes later, was really a question about the split.

The law is often also written in the rearranged form:

$$\mathbf{\Delta U = Q - W},$$

which says the same thing with the terms shuffled around: *the change in internal energy equals the heat supplied minus the work done by the gas.* Both forms are correct. A reader should be comfortable moving between the two without fuss. What must not change, whichever form is in use, is the meaning of the three symbols. **Q** is always *the heat supplied to the gas by its surroundings.* **W** is always *the work done by the gas on its surroundings.* **ΔU** is always *the change in the internal energy of the gas itself.* Reverse any of those three definitions, and the signs start to go the wrong way.

Sign convention

A sign convention is the most important piece of mental hygiene in this chapter. More students lose marks on NECTA thermodynamics questions through sign errors than through any other single mistake, and the only defence is to decide on the convention at the start and never deviate from it. The one we adopt is the standard one, locked into the form of the first law written above, and every worked example in this chapter will rehearse it explicitly before quoting a final answer.

Symbol	Positive means	Negative means	Unit
Q	heat enters the system	heat leaves the system	J
W	gas expands; gas does work on surroundings	gas is compressed; surroundings do work on gas	J
ΔU	internal energy rises; temperature rises	internal energy falls; temperature falls	J

A habit worth installing immediately: when the numbers at the end of a calculation come out with a sign that contradicts the physics of the situation, such as an answer that says the gas cooled while it was being compressed, or that the gas released heat while it was being flamed from below, the mistake is almost always in a sign, not in the formula. Verifying the sign of each quantity against physical expectation, before the final answer is stated, is not a habit to be learned in a panic the night before an exam. It is a habit to be installed right now, in the first worked example of the section, and rehearsed in every one that follows.

The law on which the first law rests

The first law of thermodynamics is not really a new law. It is the **law of conservation of energy**, applied to a system that exchanges energy with its surroundings through two particular channels: **heat** and **work**. Conservation of energy is itself one of the oldest and most secure principles in physics. It was known in its mechanical form in the eighteenth century, when it was noticed that the kinetic energy of a falling body equals the loss of its gravitational potential energy, no matter the shape of the path it falls along. What the first law adds to that older statement is the recognition that heat, which was once thought to be a separate substance called caloric that flowed invisibly from hot things to cold things, is in fact just another form of energy, and can be traded for mechanical work in a strict ratio that never changes.

The ratio is worth remembering, because it is the experimental foundation of the subject. Every 4.186 joules of mechanical work, fully converted into heat, raises the temperature of one gram of water by one degree Celsius. The figure was established by James Prescott Joule in the 1840s, in a series of patient and beautifully designed experiments that used a falling weight to turn a paddle inside an insulated tank of water. The falling weight did a measurable amount of mechanical work. The paddle stirred the water. The friction between the paddle and the water converted that mechanical work into heat. Joule measured the temperature rise of the water, weighed the weight, measured the height it fell through, and computed the ratio of mechanical work to heat produced. He repeated the experiment with different weights, different heights, different volumes of water, and even with different liquids, and he always obtained the same number. His answer has not been improved by more than a small correction in the hundred and eighty years since. The unit of energy we use today is named after him, and the first law is in some sense simply the promotion of his observation from a single experimental ratio into a universal principle: *mechanical work and heat are two forms of the same thing, and the books of energy, written in joules, must balance whenever the two are exchanged.*

Evidence for the validity of the first law

The first law has never been caught breaking. Every serious attempt to construct a counterexample has failed. The evidence for its validity comes from four independent directions, each of which would be compelling on its own, and which together leave the law on an experimental foundation as solid as anything in physics.

1. Joule's paddle-wheel experiment

Repeated over the last century and a half with ever more refined apparatus, Joule's experiment always gives the same mechanical-to-heat conversion ratio. The ratio is universal: it does not depend on the material being heated, on the mechanism

by which the mechanical work is done, or on the scale of the apparatus. This universality is exactly what one would expect from a conservation law, and exactly what one would not expect from anything else.

2. Electrical heating

When a current passes through a resistor, electrical energy is converted into heat at a rate predicted exactly by the power dissipated in the resistor. Electric kettles, of the kind sitting on Mr. Akilikubwa's desk, work on this principle, and they convert their input electrical energy into heat with an efficiency close to one hundred percent. If any joule of the input were going missing, the mismatch would show up in a careful measurement of electrical energy supplied versus thermal energy released. No such mismatch has ever been found. Billions of kettles have been boiled by this physics. None has cheated.

3. The absence of any perpetual motion machine of the first kind

A **perpetual motion machine** of the first kind *is a machine that produces more energy than it consumes, or produces energy from no source at all.* Such a machine would solve every energy problem the world has ever had, and its inventor would become the wealthiest human in history. Thousands have been proposed, from medieval overbalanced wheels to modern patents involving hidden magnets and claims of overunity generation. None has ever worked. Patent offices around the world now routinely decline to examine such proposals, because three centuries of experience have made clear that the first law will defeat them all.

4. The balance sheet of every real engine

Every engine ever operated, from the steam engines of eighteenth-century England to the gas turbines at Ubungo power station in Dar es Salaam, has consumed at least as much energy in fuel as it has produced in useful work. A careful audit of any real engine, accounting for the fuel energy burned, the waste heat rejected, and the useful work extracted, always finds that the books balance. In every case the amount of fuel energy consumed is at least as large as the useful work produced. In fact it is considerably larger, because most of the fuel energy escapes as waste heat rather than becoming useful work, but the reason that this should be so is not a matter of the first law. It is a matter of the second, which is waiting at the end of this chapter.

Limitations of the first law

For all its power and all its evidence, the first law is not the whole story. It has limitations, and being clear about them prepares the ground for the second law at the end of the chapter. Three of these limitations are worth naming explicitly, and every A-level student should be able to articulate them.

1. The first law does not predict the direction of a process

Consider a cup of hot tea on a desk in a cool classroom. Over the next hour the tea will cool, and the classroom air will warm very slightly, until the two are at the same temperature. This is the familiar direction. Now imagine running the film backwards: the cooled tea spontaneously heating itself up, while the classroom air gives up a tiny fraction of its thermal energy, until the cup of tea is once again steaming. No joule is missing in the backwards film. Every joule that left the tea in the forward process is simply flowing back into it. The first law is perfectly satisfied by the backwards film. But the backwards film never happens. Nobody has ever seen a cup of tea warm itself by extracting heat from a cooler room, and nobody ever will. The first law does not forbid this process, and that silence is its first limitation.

2. The first law does not distinguish between high-quality and low-quality forms of energy

From the first law's point of view, a joule of mechanical work and a joule of heat are interchangeable; both are just a joule. But in practice, they are not interchangeable. A joule of mechanical work can always be completely converted into a joule of heat, as happens every time someone rubs their hands together on a cold Iringa morning, or every time a brake pad dissipates kinetic energy on a long descent from the Usambara escarpment. The reverse is not true. A joule of heat can never be completely converted back into a joule of mechanical work, no matter how ingenious the machinery. Some fraction is always lost, rejected to a cooler reservoir. Heat is a lower-quality form of energy than work, because work freely becomes heat while heat does not freely become work. The first law, by counting both in joules, erases this distinction, and to that extent it under-describes the physics.

3. The first law sets no ceiling on the efficiency of energy conversion

Because the first law does not distinguish between heat and work by quality, it cannot predict how efficiently one can be converted into the other. A heat engine supplied with a thousand joules of heat might, for all the first law knows, deliver anywhere from zero to a thousand joules of useful work. Engineers designing real engines need to know the ceiling in advance, and the first law cannot tell them where it is.

All three limitations share a common theme. They concern what the first law does not predict, not what it predicts wrongly. Within its proper domain, which is the tracking of energy between the three accounts of heat, work, and internal energy, the first law has an unbroken record of two centuries. For the questions that lie outside that domain, such as which direction a process will spontaneously run in, or how efficient an engine can possibly be, a deeper principle is needed.

Enough theory. Four worked examples now follow, each short, each rehearsing the sign convention, and each chosen so that a different species of sign error gets its audition. We begin with the tamest.

BINDER Example 1

A gas absorbs 500J of heat, and during the process its internal energy increases by 200J. Find the work done by the gas, and describe in physical terms what the gas is doing.

Solution

Apply the first law in its standard form, with every sign decided by the physics of the situation before any numbers are substituted in. The gas absorbs heat, so heat enters the system, which in our sign convention means Q is positive: Q = +500J. The internal energy of the gas increases, which means the gas becomes hotter, consistent with heat going in and the temperature rising, so ΔU is positive: ΔU = +200J. The work done by the gas is what is being asked for, and the first law supplies it directly:

$$W = Q - \Delta U = 500J - 200J = 300J$$

The sign of W is positive, which means the gas has done 300J of work on its surroundings.

Making Sense of the Answer: *Of the 500J that entered the gas as heat, 200J stayed inside it as raised internal energy, in the form of hotter, faster-moving molecules, and 300J left again as mechanical work done on the surroundings. The two destinations add up to 500J exactly, as the first law requires. In physical terms, the gas is expanding against its surroundings while being heated from below. One might picture a gas in a cylinder with a movable piston, warmed by a flame, pushing the piston upward while its temperature also rises. Every joule is accounted for.*

Think Like a Physicist: *The first law has told us how the books balance, but it has said nothing about the mechanism. Is the gas lifting a piston? Pushing back the atmosphere? Turning a turbine blade? The equation is silent on the question. It guarantees only that the total is 300J of work and 200J of retained heat, whatever the arrangement happens to be. This is both the power of the law, since it is universal and applies equally to any closed system, and its deliberate limitation: the law does not care about the particular piston. A physicist who wants the particular piston must bring in a separate piece of physics, which in this chapter will be the PV diagram of the next section.*

BINDER Example 2

A gas is compressed quickly by an external force, which does 300J of work on the gas. During the compression, 120J of heat leaks out of the gas into the surroundings. Find the change in internal energy of the gas.

Solution

This example is designed to rehearse every sign in the convention, and it rewards readers who take each symbol in turn, verifying the sign against the physical situation, before touching the equation at all.

The external force does work on the gas, which means the gas is being compressed, not expanding. In our sign convention, which defines W as the work done by the gas on its surroundings, the work done by the gas is the negative of the work done on the gas. So W = −300J. The minus sign is doing real work here. A reader who plugs +300J into the formula will obtain a final answer with the wrong sign, and will report that the gas cooled down when in fact it has warmed up. Heat leaks out of the gas to the surroundings, which in the convention means Q is negative, so Q = −120J. Now, and only now, apply the first law in its rearranged form:

$\Delta U = Q - W = (-120J) - (-300J) = -120J + 300J = +180J.$

Making Sense of the Answer: *The change in internal energy is positive. The gas has become hotter, which is exactly what one expects when a gas is rapidly compressed. The compression added 300J of work to the internal budget, and the heat leakage removed 120J, leaving a net gain of 180J. This is the physics of a bicycle pump barrel getting warm as air is rapidly pushed through it, which is precisely what Kipanga's pump was doing in the opening pages of chapter 4. It is also the physics of a diesel engine cylinder heating up fast enough to ignite fuel without a spark, which is one of the applications we will meet later.*

Think Like a Physicist: *The two minus signs that appeared in this problem are not a quirk of algebra; they are the physics. Each minus sign is a separate decision about which direction the energy is flowing. One minus sign records that work is being done on the gas rather than by it. The other records that heat is leaving the gas rather than entering it. A student who treats signs as a mechanical recipe, plugging numbers in hopefully and checking at the end, will fail on this class of problem. A student who treats each sign as answering the question 'which way?' will succeed. The habit to build is not 'put a minus sign in', but 'decide, before I write anything down, whether this quantity is going in or coming out'. Build that habit here, on a simple numerical problem, because it is the habit that carries you through the adiabatic process in a later section.*

REAL Example 3

On a cold Iringa morning, a student waiting outside the assembly hall rubs her palms together vigorously. Within a few seconds her palms become noticeably warm, even though no flame, no stove, and no electrical heater is anywhere near her.

(a) Using the first law of thermodynamics, explain the physical mechanism.
(b) A few seconds after she stops rubbing, her palms begin to cool back down to the temperature of the morning air. What form does the first law take during that cooling phase?

Solution

(a) The first law of thermodynamics applies to any closed system, not only to a gas inside a container. For this example, take the two palms together as the system. The surroundings are the rest of the student's body, the cold Iringa air, the inside of her school uniform, and everything else beyond the outer surface of her skin.

While she is rubbing, the muscles of her arms and shoulders are doing mechanical work on her palms through the friction between them. Friction converts mechanical work into heat inside the material where the friction occurs, which in this case is the skin of the palms themselves. In the language of the first law, the palms are receiving energy as work rather than as heat. The system is a solid, not a gas, so there is no significant expansion, which means no energy is leaving again as expansion work on the surroundings. All of the incoming work stays in the palms, raising their internal energy. Hotter molecules in the palms mean a higher temperature, and the palms feel warm. The balance during the rubbing phase is therefore

$$\Delta U = -W \quad \text{(palms gain internal energy from work done on them),}$$

with the heat term Q approximately zero over the few seconds of rubbing, because the palms have not had time to lose any appreciable heat to the cold air around them.

(b) When she stops rubbing, the external work stops. The palms are now warmer than the surrounding air, so heat begins to flow outward from palms to air, exactly as chapter 2 described: heat flows spontaneously from the hotter body to the cooler one. In the first law, the work term is now zero, and the heat term is negative because heat is leaving the system. The balance becomes

$$\Delta U = Q \quad \text{(palms lose internal energy as heat escapes to the air),}$$

with Q negative. The internal energy of the palms decreases, which is the physical sensation of her palms cooling back down to the morning temperature.

Making Sense of the Answer: *The same first law, written in exactly the same form, describes both halves of the situation. What changes between the two halves is not the law itself but which of the two channels, work or heat, is active. During the rubbing, the work channel is carrying energy in and the heat channel is quiet. After the rubbing, the work channel is quiet and the heat channel is carrying energy out. The student's palms behave as a thermodynamic system exactly like a kettle, an engine cylinder, or the air inside a sealed tyre. The same accounting runs in all four cases.*

Think Like a Physicist: *Notice what this example has done for us. It has taken the first law, which was stated originally for a gas, and applied it to a pair of human palms. The palms are solid tissue, not gas; there is no piston, no container, no pressure; and yet the first law applies without modification. This is the lesson. The first law is not a law about gases. It is a law about closed systems. Any closed system will do: a gas, a liquid, a solid, a pair of hands, a rock warming in the sun, the earth itself radiating to space. The skill in applying the first law to an unfamiliar situation is never mathematical. It is always conceptual, and it always comes down to three decisions. What is the system. What counts as heat. What counts as work. Once those three decisions are made, the arithmetic writes itself.*

HOT Example 4

One gram of water at 100°C is vapourised at atmospheric pressure into 1671cm^3 of steam at the same temperature. The specific latent heat of vapourisation of water at 100°C is $2.26\times10^6\text{Jkg}^{-1}$, and atmospheric pressure is $1.013\times10^5\text{Pa}$. Calculate:

(a) the heat supplied to the water during the vapourisation,
(b) the external work done by the expanding water against the atmosphere, and
(c) the increase in the internal energy of the water as it becomes steam.

Solution

This example is a direct bridge between chapter 2 and chapter 5, and the bridge is worth making explicit as we go. The heat supplied comes from chapter 2's latent-heat formula, because the process is a phase change at constant temperature rather than a simple heating of a fixed-phase sample. The external work is a new chapter 5 quantity, computed from the volume change against a constant external pressure. The internal energy change follows directly from the first law, once the other two are in hand. We take the three parts in order.

(a) The water is vapourising at constant temperature, so there is no temperature rise at all, and the heating formulae that involve a temperature change do not apply. What applies instead is chapter 2's latent-heat formula, Q = mL, where m is the mass of water vapourising and L is the specific latent heat of vapourisation. Substituting:

$$Q = (1.0\times10^{-3}\text{kg})(2.26\times10^6\text{Jkg}^{-1}) = 2.26\times10^3\text{J} = 2260\text{J}$$

The sign is positive, because heat is being supplied to the water from the flame below.
(b) The water begins as a liquid with a volume of about 1cm^3 (with $\rho_{\text{water}} = 1\text{gcm}^{-3}$, 1g of liquid water has a volume of 1cm^3 at this temperature) and ends as steam with a volume of 1671cm^3. The change in volume is therefore

$$\Delta V = 1671\text{cm}^3 - 1\text{cm}^3 = 1670\text{cm}^3 = 1.670\times10^{-3}\text{m}^3$$

Throughout the expansion, the system is pushing against the surrounding atmosphere at a constant pressure of 1.013×10^5Pa, so the work done by the expanding water on the atmosphere is:

$$W = P\Delta V = (1.013\times10^5\text{Pa})(1.670\times10^{-3}\text{m}^3) = 169\text{J}$$

The sign is positive, because the system is expanding and therefore doing work on its surroundings.

(c) Change in internal energy. With Q and W in hand, the first law supplies ΔU directly:

$$\Delta U = Q - W = 2260\text{J} - 169\text{J} = 2091\text{J}$$

Making Sense of the Answer: *Three numbers now sit side by side, and the relationship between them is the whole point of the example. Of the 2260J supplied to vapourise one gram of water, only 169J went into mechanical work against the atmosphere. That is about seven percent of the total. The remaining 2091J, which is roughly ninety-three percent, stayed inside the water as an increase in internal energy, even though the temperature did not rise by a single degree. The vapour molecules are at exactly the same temperature as the liquid water they came from, but they are no longer held close to each other by the intermolecular forces that existed in the liquid. Separating those molecules from each other, against their own mutual attraction, takes energy, and that energy is now locked inside the vapour as an increase in its internal energy. This is the physical content of latent heat. Most of the energy supplied during boiling goes into breaking molecular bonds, and only a small fraction goes into pushing back the atmosphere.*

Think Like a Physicist: *Notice what the first law has just done. It has stitched together three quantities that at first sight belong to three different chapters. The heat supplied Q came from chapter 2's latent-heat formula. The work done W came from chapter 5's picture of expansion against atmospheric pressure. And the increase in internal energy ΔU reached back to chapter 4's molecular account of what internal energy actually is, namely the kinetic energy of molecular translation, supplemented by rotational and vibrational modes, but now also including the energy stored in separating molecules against the intermolecular forces that were holding them together. The first law does not care which chapter supplied each of its three terms. It cares only that the three terms balance, which they do, to the last joule.*

Four worked examples now lie behind us, and the first law has emerged from each of them with its books still balancing and its dignity intact. What has not yet been cracked open is the work term itself. W has been a single number so far, pulled from either algebra (Example 1) or a product PΔV at fixed pressure (Example 4). Both cases are special; in a general thermodynamic process, the pressure declines to hold still while the gas expands, and W must be computed by an integral rather than a single multiplication. The integral, graphed, turns into an area; and the diagram on which that area lives is the most useful picture in this chapter. That picture is called the PV diagram, and the next section introduces it.

WORK DONE BY A GAS

The previous section ended with a promise. The work term W was going to open up, and the tool for opening it was going to be a diagram. That is now overdue. Every worked example so far has treated W as a single number, pulled either from the first law by arithmetic or from a product PΔV at fixed pressure. Both of those are special cases. The general rule, for any process at all, is that the work done by a gas is not a number but an integral, and the integral has a picture, and the picture carries the physics. That picture is called the **PV** diagram, and the next few pages live on it.

Before the diagram, we need the integral. Before the integral, we need the piece of physics the integral integrates. That piece, in turn, needs only one picture to see clearly, and the picture is this one.

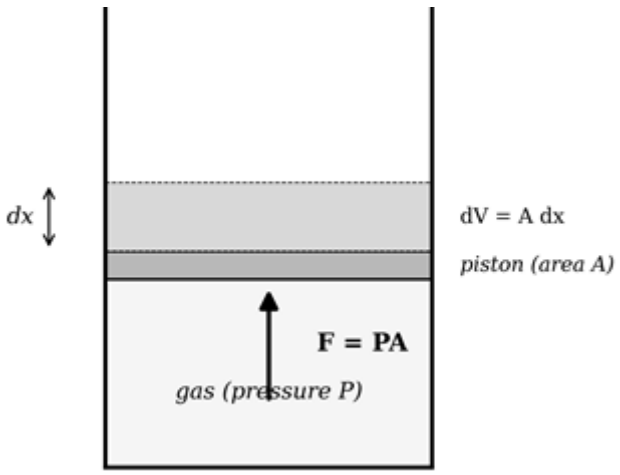


Figure: ***A gas trapped inside a cylinder with a movable piston of cross-sectional area A.*** *The gas exerts an upward force F = PA on the piston, where P is the gas pressure. When the piston slides outward by a small distance dx, the volume swept out (shown shaded) is dV = Adx, and the small quantity of work done by the gas on the piston is dW = Fdx = PdV. This tiny identity is the foundation of everything that follows.*

From force on a piston to work done by a gas

Picture the cylinder in the figure. A gas at pressure P fills the lower half; a piston of cross-sectional area A closes the top. The gas molecules, in the course of their chaotic motion from Chapter 4, are constantly colliding with the underside of the piston. Each collision transfers a tiny amount of momentum upward. Averaged over all molecules and all collisions, these transfers add up to a steady net force on the piston, directed upward, and equal, by the very definition of pressure, to the pressure multiplied by the area on which it acts:

$$F = PA$$

This is not new physics. It is the pressure formula from Chapter 4, applied to one particular wall of the container. What is new is what we do with it next.

Let the piston slide outward by a small distance dx. Nothing mysterious happens in this motion; the gas pushes, the piston moves, work is done by the pusher on the pushed. The small amount of work done by the gas on the piston in this small motion is the force multiplied by the distance:

$$dW = Fdx = PAdx$$

Now comes the observation that turns this piece of mechanics into a piece of thermodynamics. The volume swept out by the piston as it moves through dx is Adx, and this is exactly the change in the volume of the gas:

$$dV = Adx$$

Substituting Adx = dV into the expression for dW gives the result of this section in a single line:

$$\mathbf{dW = PdV}$$

The small quantity of work done by a gas, when its volume changes by a small amount dV, is the pressure multiplied by that volume change. The identity holds for a piston, for the hovering lid of a kettle, for the expanding gases in a car engine, and for any other situation in which a gas expands against some movable boundary. The boundary does not have to be a flat piston; the formula is a local statement about an infinitesimal volume change, and any shape of boundary can be broken down into many small flat pieces.

For a finite change in volume, from some initial V_1 to some final V_2, the total work done is the sum of many such small contributions. In the language of calculus, this sum is an integral:

$$\mathbf{W = \int_{V_1}^{V_2} PdV}$$

This is the central formula of the section. It says that the **work done** by a gas during any process, whether slow or fast, whether the pressure stays fixed or changes wildly, whether the gas is ideal or real, *is the integral of the pressure with respect to the volume, taken along the path of the process.* A student who carries this formula into the rest of the chapter can derive everything else.

The integral is the area under the curve

An integral of the form ∫PdV has an immediate graphical meaning that is worth pausing on, because it is the reason the PV diagram is worth drawing in the first place. If we plot P on the vertical axis against V on the horizontal axis, then during any thermodynamic process the state of the gas traces out a curve on this diagram. The quantity PdV, for a thin vertical strip of width dV under the curve, is exactly the area of that strip: height P, width dV, area P dV. The integral, which sums all these strips from V_1 to V_2, is therefore the total area enclosed between the curve, the V-axis, and the two vertical lines $V = V_1$ and $V = V_2$. In a single phrase:

$$\mathbf{W = \text{area under the PV curve}}$$

This is how the PV diagram does physics. It makes the work done by the gas visible as an area on a page. Two processes whose curves enclose the same area do the same amount of work; two processes whose curves enclose different areas do different amounts of work; and the shape of the curve matters as much as its endpoints. ***Foundations of Physics*** taught us that the area under a force-displacement graph equals the work done, which is exactly the same story as the one told here, translated from mechanical units into thermodynamic ones.

Checking the sign against the first-law convention

The sign convention from the previous section required W positive for expansion and negative for compression. The integral gets this right automatically, provided the limits are written in the order the process actually runs. When the gas expands, the final volume V_2 exceeds the initial volume V_1, the integration runs from a smaller number to a larger one, and since the pressure is positive throughout, the integral is positive. $W > 0$: the gas has done work on its surroundings. When the gas is compressed, V_2 is less than V_1, the integration runs from a larger number to a smaller one, the integral comes out negative, and $W < 0$: the surroundings have done work on the gas. When the volume does not change at all, $V_2 = V_1$, the integral vanishes, and $W = 0$. That last case is the isochoric process, which gets its own subsection soon.

The convention locked at the start of the previous section survives intact. No new rules are needed; the integral and the sign convention already agree.

Work is path-dependent

Here is the observation that makes this subject interesting, and that the rest of the chapter will build on. The integral $\int PdV$ depends on the shape of the curve between V_1 and V_2, and not merely on the endpoints. Two processes connecting the same initial state $A = (V_1, P_1)$ to the same final state $B = (V_2, P_2)$ can trace out very different curves, enclose very different areas, and therefore deliver very different amounts of work. This is true even though both processes start at the same place and end at the same place. Work is a **path function**, not a **state function**. Its value cannot be recovered from the endpoints alone; the whole history of the process matters.

The figure below makes the point graphically. Three different processes carry the gas from the same point A in the upper left to the same point B in the lower right, each along a differently shaped curve, and each enclosing a visibly different area beneath itself. The three areas are the three values of W.

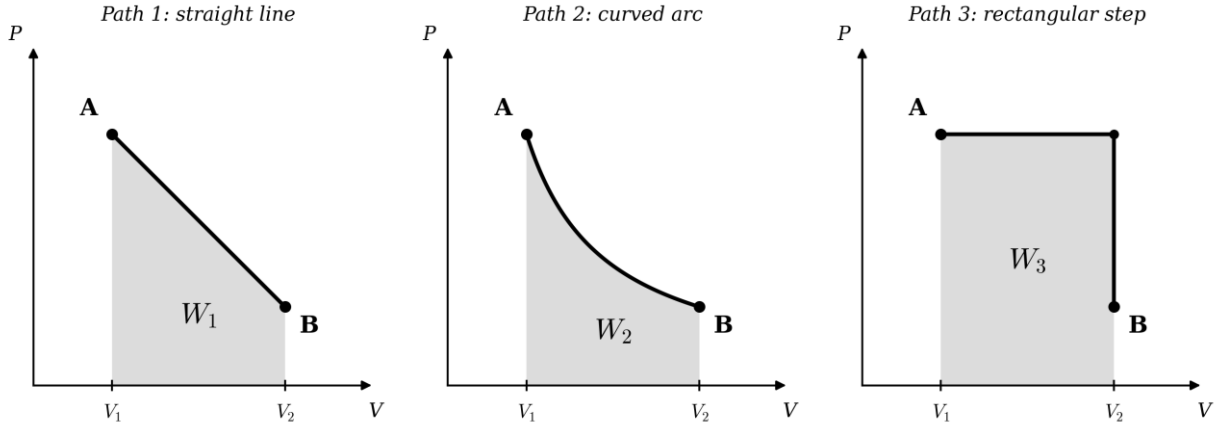


Figure: ***Three thermodynamic processes connecting the same initial state A to the same final state B, along three different paths.*** *The shaded area beneath each path is the work W done by the gas during that process. The three areas are visibly unequal: W_1 (straight line), W_2 (curved arc, lying below the line), and W_3 (rectangular step, lying above the line). Same endpoints, three different answers. Work is path-dependent.*

The curved arc of Path 2 lies everywhere below the straight line of Path 1, so the area under Path 2 is smaller than the area under Path 1, and $W_2 < W_1$. The rectangular step of Path 3 goes first along the top of the box at constant high pressure, then straight down at constant volume; the area under this step is the full rectangle of width $V_2 - V_1$ and height P_1, which is everywhere above the straight line, so $W_3 > W_1$. All three processes carry the gas from the same A to the same B, yet all three do different amounts of work against the surroundings.

State functions and path functions

The path-dependence of W is worth contrasting with another quantity that behaves very differently. The state variables P, V, T, and n are, as their name suggests, functions of state: each has a definite value at every equilibrium state, and the change in any of them between two states depends only on the two states, not on how the gas got from one to the other. If the gas ends at $V = V_2$, the volume is V_2 regardless of whether the gas reached V_2 by a straight line, a curved arc, or a rectangular step. The same is true of the pressure, the temperature, and the amount of substance. It is also true of the internal energy U, which Chapter 4 showed depends only on the temperature for an ideal gas, and therefore only on the state, not the path. The change ΔU between two states is the same no matter which process connected them.

Work W does not behave this way. Neither does heat Q. Both are path functions, not state functions. There is no such thing as "the work of state B" or "the heat of state B"; there is only the work done and the heat supplied during a particular process, and both depend on the path that process took through state space. ***This is why the symbols ΔW and ΔQ are wrong on principle and should never be written:*** there is no W of state A and no W of state B to take the difference of. The correct symbols are simply W and Q, each referring to the full accumulation over the path. The first law, $Q = \Delta U + W$, has a delta **only** on U, precisely because only U is a state function of the three.

This distinction is not a nicety. It is the reason the four thermodynamic processes about to be met, namely isochoric, isobaric, isothermal, and adiabatic, give four genuinely different answers for the work done, even when their endpoints are arranged to coincide. Each process traces a different path. Each path encloses a different area. Each area is the work done. The physics is in the curves.

This is enough theory for one section. The formula $dW = PdV$, the integral $W = \int PdV$, the area-under-curve reading, and the path-dependence that results are the four pieces that every subsequent calculation uses. Four worked examples now put them to work on increasingly serious geometry: a constant-pressure slab (the rectangle under a horizontal line), a rectangular step

(two perpendicular legs), a comparison between two different paths connecting the same endpoints, and a Tanzanian football at the Miono school field.

BINDER Example 5

A gas in a cylinder with a loose, frictionless piston is heated slowly so that its pressure remains constant at 2.0×10^5Pa throughout the process. Its volume expands from 3.0L to 8.0L. Calculate the work done by the gas during the expansion.

Solution

The pressure is constant throughout the process, so it comes out of the integral as a common factor, leaving only the integration over the volume change:

$$W = \int_{V_1}^{V_2} P\,dV = P\int_{V_1}^{V_2} dV = P(V_2 - V_1) = P\Delta V$$

Before substituting, the volumes must be converted from litres to cubic metres (the SI unit), because the pressure has been given in pascals, and the product PΔV will only come out in joules if every factor carries SI units:

$$V_1 = 3.0\text{L} = 3.0\times10^{-3}\text{m}^3,\ V_2 = 8.0\text{L} = 8.0\times10^{-3}\text{m}^3$$

The volume change is therefore:

$$\Delta V = V_2 - V_1 = 8.0\times10^{-3}\text{m}^3 - 3.0\times10^{-3}\text{m}^3 = 5.0\times10^{-3}\text{m}^3$$

Substituting into the formula for the work:

$$W = P\Delta V = (2.0\times10^5\text{Pa})(5.0\times10^{-3}\text{m}^3) = 1.0\times10^3\text{J} = 1000\text{J}$$

The work done by the gas is 1000J.

Making Sense of the Answer: *Geometrically, the PV diagram for this process is a horizontal line at $P = 2.0\times10^5$Pa, running from V = 3.0L to V = 8.0L. The work done by the gas is the area of the rectangle beneath this line: height P, width ΔV, area PΔV. This is the simplest possible PV-diagram calculation, because the shape is a rectangle and the area is a single multiplication. All 1000J of work flowed from the gas to whatever mechanism was holding the piston in place, whether a stone on top, the atmosphere, a spring, or a thoughtful lab technician's hand.*

Think Like a Physicist: *Constant-pressure processes are the gentle giants of thermodynamics. Their PV curve is a horizontal line, their area is a rectangle, and their work is a single multiplication. Whenever a problem says "**the pressure remains constant**" or "**isobaric**", reach for W = PΔV without hesitation. And always, always convert the volume to cubic metres before multiplying by a pressure in pascals. Leaving V in litres gives an answer in Pa·L, which is smaller than a joule by a factor of 1000.*

BINDER Example 6

A gas starts at state ($V_1 = 2.0$L, $P_1 = 1.0\times10^5$Pa) and undergoes a two-step process on the PV diagram. In step one, it is heated at constant volume until its pressure has risen to $P_2 = 3.0\times10^5$Pa. In step two, at this new constant pressure, its volume expands to $V_3 = 5.0$L. Calculate the total work done by the gas during the two-step process, and identify which step contributed what.

Solution

The two steps are separate integrals on the PV diagram. Step one is a vertical segment (constant volume), and step two is a horizontal segment (constant pressure). Each is computed separately and the two results are added.

Step one (isochoric heating):

The volume is constant at V = 2.0L throughout the step, which means dV = 0 at every point along it. Therefore:

$$W_1 = \int P\,dV = 0$$

No volume change means no work done by the gas, no matter how much the pressure happens to rise during the step. This is a slightly counter-intuitive fact for students meeting it for the first time, so it is worth reading twice. The gas is being heated vigorously; its molecules are hitting the walls faster and harder; the pressure is tripling. And yet the gas is doing no work, because nothing is moving. Work, in the mechanical sense, requires a force to be applied through a distance, and here the distance is zero.

Step two (isobaric expansion):

The pressure is constant at $P = 3.0\times10^5$Pa while the volume expands from 2.0L to 5.0L. Converting to SI units:

$$V_2 = 2.0\text{L} = 2.0\times10^{-3}\text{m}^3,\ V_3 = 5.0\text{L} = 5.0\times10^{-3}\text{m}^3,\ \Delta V = 3.0\times10^{-3}\text{m}^3.$$

The work done during step two is:

$$W_2 = P\Delta V = (3.0\times10^5\text{Pa})(3.0\times10^{-3}\text{m}^3) = 900\text{J}$$

The total work done by the gas during the two-step process:

$$W = W_1 + W_2 = 0 + 900\text{J} = 900\text{J}$$

Making Sense of the Answer: *On the PV diagram, the two-step path looks like an L lying on its side: a vertical leg going up (isochoric heating, no width, no area) joined at the top to a horizontal leg going right (isobaric expansion, a genuine rectangle). Only the horizontal leg contributes to the work, and it contributes its full rectangular area of 900J. The vertical*

leg, despite involving a dramatic tripling of the pressure, contributes nothing at all, because its width on the V-axis is zero. All of the work comes from the second step.

Think Like a Physicist: *This example encodes two facts that apply to every multi-leg problem. First, vertical segments on a PV diagram do zero work, always, regardless of how the pressure changes along them. Second, horizontal segments do PΔV of work, equal to the rectangular area beneath them. For any staircase path made of purely horizontal and purely vertical pieces, the total work is simply the sum of the horizontal rectangles. This observation is a fast way to estimate the work for any rectangular-step process, and it foreshadows the next few sections, where the isobaric process (all horizontal) and the isochoric process (all vertical) appear as two of the four thermodynamic workhorses.*

REAL Example 7

Kipanga is inflating a football with a hand pump at the Miono school field. On his first attempt, in no hurry, he pumps slowly and pauses between strokes. On his second attempt, running late for a match against the neighbouring school, he pumps quickly and does not pause. Both times, the football ends at the same final pressure. Kipanga notices that the pump barrel gets noticeably hotter after the fast pumping than after the slow pumping. Explain why, using the first law of thermodynamics.

Solution

Take the air inside the pump barrel as the system, and consider what happens during one compression stroke. The handle pushes down; the air compresses; the pressure inside the barrel rises until the one-way valve opens and the air moves into the football. Whether Kipanga pumps slowly or quickly, the air starts at atmospheric pressure and ends at football pressure. The two endpoints agree.

What differs between the two strokes is the amount of heat the air exchanges with the barrel walls during the stroke itself. Compressing a gas always tries to heat it up (molecules get squeezed into less space, they hit the walls harder, the temperature rises). In a slow stroke, there is time for that extra thermal energy to leak out to the barrel walls and from there to the surrounding air, so the air inside the barrel stays close to room temperature throughout. In a fast stroke, there is no time for the heat to escape; it stays trapped inside the air, so the air gets genuinely hot by the end of the stroke.

Now apply the first law, $Q = \Delta U + W$, to the air in the barrel for one stroke. In the slow stroke, ΔU is small (the temperature hardly rises), Q is negative (heat leaves the air), and W is modest. In the fast stroke, ΔU is large (the air finishes much hotter), Q is nearly zero (no time to lose heat), and W is larger. By the first law, the extra work done by Kipanga's arm in the fast stroke has nowhere to go except into ΔU, so the air finishes hotter. That hot air, once it passes through the valve, dumps its heat onto the valve and the upper barrel walls. Stroke after stroke, those walls accumulate heat faster than they can shed it to the surroundings. The barrel gets warm. The slow stroke deposits very little heat per stroke, and the barrel has plenty of time between strokes to cool, so it stays near ambient.

Making Sense of the Answer: *Both pumpings end with the same air in the same football at the same final pressure. The endpoints agree. But the amount of work Kipanga's arm has done on the air, and the amount of heat the air has shed to the barrel, differ between the two paths. This is path-dependence of W and Q in its most everyday form: a football pump whose temperature reports on how fast you were pumping.*

Think Like a Physicist: *The barrel's temperature is a readout, not a cause. Kipanga's arm was the real source of the heat. The first law is what converts his muscular work, via the intermediate step of compressed air, into the warmth of the barrel walls. When a thermodynamic system behaves differently at different speeds, the first question is always whether the path is different, and the second is how ΔU, Q, and W are rearranging themselves along that path. The endpoints being the same does not settle anything; the path is where the physics lives.*

HOT Example 8

A gas is taken from state A, where $P_A = 3.0\times10^5$Pa and $V_A = 1.0\times10^{-3}$m³, to state B, where $P_B = 1.0\times10^5$Pa and $V_B = 3.0\times10^{-3}$m³, along two different paths on the PV diagram. Path 1 is a straight line directly from A to B. Path 2 consists of an isobaric expansion at $P = 3.0\times10^5$Pa from V_A to V_B, followed by an isochoric drop at $V = V_B$ from P_A down to P_B. Compute the work done by the gas along each path.

Solution

The initial and final states A and B are fixed. Only the shape of the path between them differs. We compute each path separately by finding the area beneath its PV curve. The two paths are shown in the figure below; the coordinates of A and B are given in the display units of the axes, and the shaded region in each panel is the area whose value we are about to compute.

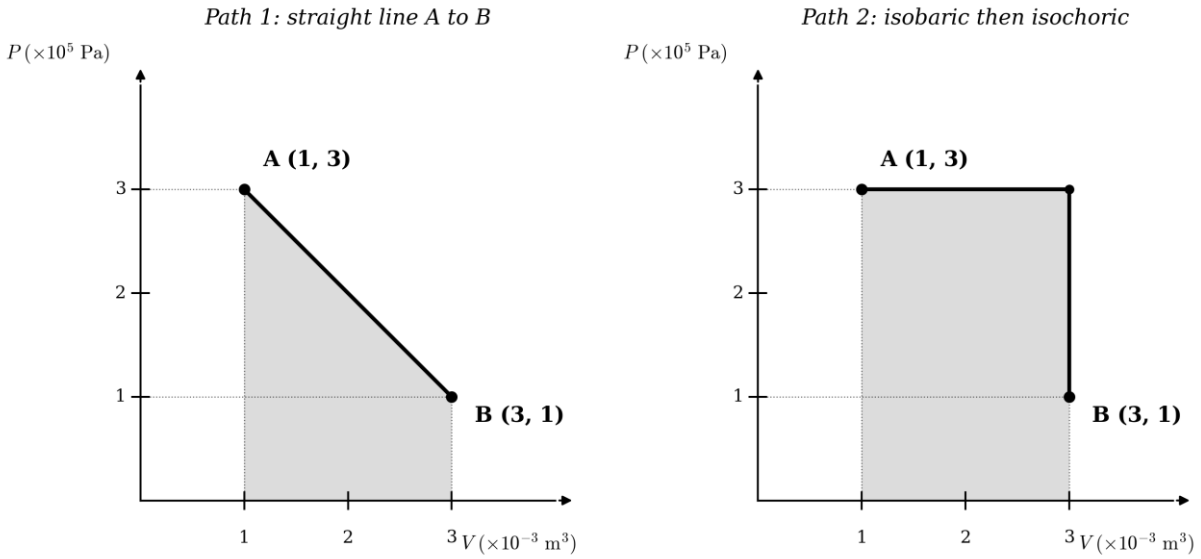


Figure: ***The same initial state A and final state B, connected by two different paths****. The coordinate pair in brackets next to each point gives (V, P) in the display units of the axes. On the left, a straight line from A to B encloses a trapezoidal area beneath it. On the right, an isobaric expansion at the upper pressure (horizontal leg) followed by an isochoric pressure drop at V_B (vertical leg) encloses the full rectangle beneath the horizontal leg. The rectangle is visibly larger than the trapezoid, and the calculation that follows turns that visual difference into a precise number of joules.*

Path 1: the straight line

This encloses a trapezoidal area, so:

$$\mathbf{W_1} = \frac{P_A - P_B}{2} \times (V_B - V_A) + P_B(V_B - V_A) = \frac{\mathbf{P_A + P_B}}{\mathbf{2}} \times (\mathbf{V_B - V_A})$$

Substituting:

$$W_1 = \frac{3.0 + 1.0}{2} \times 10^5 \text{Pa} \times (3.0 - 1.0) \times 10^{-3}\text{m}^3 = 400\text{J}$$

The work done along Path 1 is 400J.

Path 2: isobaric expansion followed by isochoric drop

This encloses a rectangular area, so:

$$W_2 = P_A(V_B - V_A) = (3.0 \times 10^5\text{Pa})(3.0 - 1.0) \times 10^{-3}\text{m}^3 = 600\text{J}$$

The work done along Path 2 is 600J.

Making Sense of the Answer: *Same A, same B, two different values for W. The difference is real and fully quantitative: Path 2 does 200 joules more work than Path 1. Why? Path 2 stays at the high pressure* P_A *throughout the volume expansion and then drops the pressure only after the volume has finished changing. It captures the full high-pressure rectangle as its work. Path 1 lets the pressure drop gradually during the expansion, so its average pressure during the volume change is smaller, and the trapezoid beneath it encloses less area than the rectangle. The difference between the two areas is the difference between the two values of W. The physics is visible on the diagram; the diagram agrees with the arithmetic; and the arithmetic agrees with the first law, as we are about to verify.*

Think Like a Physicist: *Here is the bonus consequence of this example, and it is the reason this problem is classified as HOT. Because the endpoints A and B are identical in the two processes, the change in internal energy ΔU must also be identical in both. U is a state function; it cares only about the endpoints. Write the first law for each path:* $Q_1 = \Delta U + W_1$ *and* $Q_2 = \Delta U + W_2$*. Subtract the two equations:* $\mathbf{Q_2 - Q_1 = W_2 - W_1}$ *= 200J. The second path requires 200 joules more heat than the first path, because it does 200 joules more work, and the internal energy is the same in both cases. This is how Q and W share the path-dependence burden. Whatever extra work is done by the gas must be fed by extra heat from the surroundings, because ΔU is already pinned down by the endpoints.*

Work is now in hand. The next section takes up the other half of the first law, the heat Q, and in particular how much heat is needed to raise the temperature of a gas by a given amount.

MOLAR HEAT CAPACITIES OF AN IDEAL GAS

The previous section opened the work term W. This section opens the heat term Q. The question to be answered, in a single sentence, is this: *how much heat is required to raise the temperature of a gas by a given amount?* Chapter 2 answered the corresponding question for solids and liquids using a single specific heat capacity, because solids and liquids hardly expand when heated, and so the distinction between heating at constant volume and heating at constant pressure made no practical

difference for them. Gases are different. A gas heated at constant volume needs less heat than the same gas heated at constant pressure, for the same temperature rise, because in the second case some of the heat is being spent on expansion work against whatever is holding the pressure fixed. Two different setups, two different answers, two different heat capacities. The picture that makes this vivid is the following.

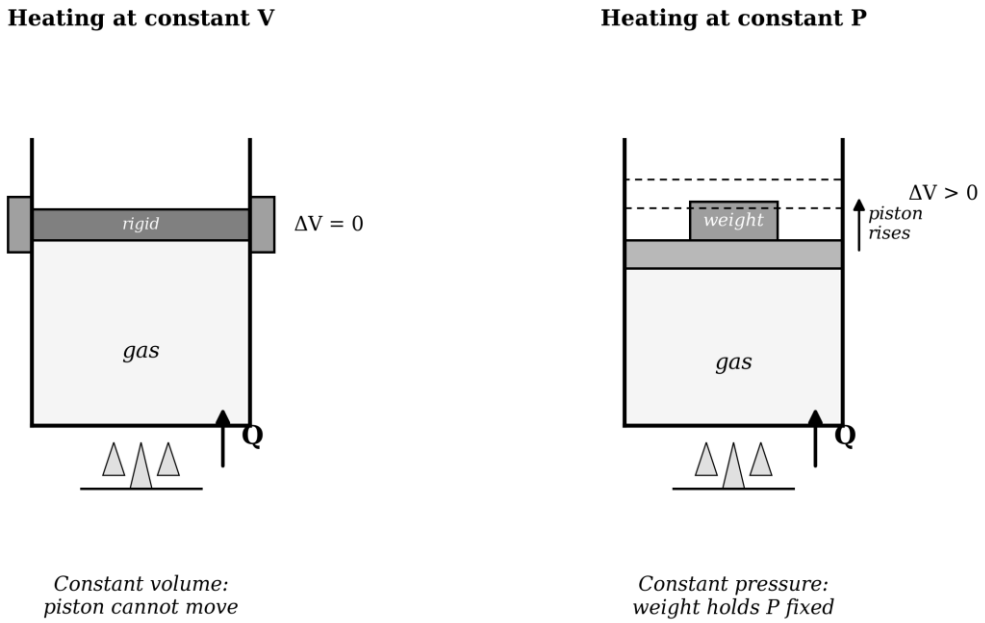


Figure: ***Two identical samples of the same gas, in two identical cylinders, are heated from below by identical flames delivering the same quantity of heat Q.*** *On the left, the piston is clamped rigid, so the gas cannot expand: the volume is held constant, and all of the incoming heat goes into raising the internal energy of the gas, which raises its temperature. On the right, the piston is free to slide, with a fixed weight on top keeping the pressure constant, so as the gas heats up it expands, lifts the weight, and does work against gravity. Some of the incoming heat becomes internal energy; the rest leaves as work. For the same Q, the right-hand gas ends at a lower temperature than the left-hand gas. Or, equivalently, to reach the same temperature rise, the right-hand gas requires more heat than the left-hand one. The difference is exactly the expansion work.*

The physics the figure encodes is the punchline of the whole section: a gas has not one but two heat capacities, one for each of the two arrangements. The left arrangement gives the **molar heat capacity at constant volume**, written $\mathbf{C_v}$. The right arrangement gives the **molar heat capacity at constant pressure**, written $\mathbf{C_p}$. Both have units of $Jmol^{-1}\ K^{-1}$, and both refer to the heat needed per mole of gas, per degree of temperature rise. The rest of the section derives each one from the molecular picture of Chapter 4, connects them through the first law, and lands on the famous relation between them: **Mayer's equation**.

Heat capacity at constant volume

Consider the left-hand cylinder in the figure. The piston is rigid, so the gas cannot expand, meaning $\Delta V = 0$ during the heating. By the result from the previous section, $W = 0$ for a process at constant volume. The first law then simplifies:

$$Q = \Delta U + W = \Delta U + 0 = \Delta U$$

Every joule of heat that flows into the gas goes into raising its internal energy. This is a very clean setup, because the heat Q is literally the change in U, nothing lost to anything else.

Now invoke the result from Chapter 4 that expressed the internal energy of an ideal gas in terms of its temperature. For a monatomic ideal gas with **n** moles and absolute temperature **T**, the **equipartition theorem** gave:

$$U = \frac{3}{2} \times nRT$$

The change in internal energy for a small temperature change ΔT is therefore:

$$\Delta U = \frac{3}{2} \times nR\Delta T$$

Combining this with $Q = \Delta U$:

$$\mathbf{Q = \frac{3}{2} \times nR\Delta T} \quad (\text{monatomic, constant V}).$$

The **molar heat capacity at constant volume** is defined as *the heat required per mole per degree temperature rise at constant volume:*

$$C_v = \frac{Q}{n\Delta T} \quad \text{at constant V}$$

Substituting the monatomic expression for Q:

$$C_v = \frac{\frac{3}{2} \times nR\Delta T}{n\Delta T}$$

$$\mathbf{C_v = \left(\frac{3}{2}\right)R} \quad \text{(monatomic ideal gas).}$$

For a diatomic gas at moderate temperatures, Chapter 4 gave five active degrees of freedom (three translational plus two rotational), so U = (5/2)nRT and the molar heat capacity at constant volume is:

$$\mathbf{C_v = \left(\frac{5}{2}\right)R} \quad \text{(diatomic ideal gas).}$$

In general, for a gas with f active degrees of freedom per molecule, equipartition gives U = (f/2)nRT and therefore:

$$\mathbf{C_v = \left(\frac{f}{2}\right)R = f \times \frac{R}{2}}$$

The molar heat capacity at constant volume is, at this level of theory, a direct measurement of *how many degrees of freedom each molecule of the gas has, multiplied by the universal constant R/2*. This is one of the most beautiful connections in the whole of classical physics: a macroscopic, thermometer-and-stopwatch quantity like C_v counts the microscopic degrees of freedom of the molecules.

Numerically, with R = 8.314 $Jmol^{-1}\,K^{-1}$, the predictions are C_v = 12.5 J $mol^{-1}\,K^{-1}$ for a monatomic gas and C_v = 20.8 J $mol^{-1}\,K^{-1}$ for a diatomic gas at room temperature. The measured values for helium, neon, and argon (monatomic) cluster tightly around 12.5. The measured values for nitrogen, oxygen, hydrogen, and carbon monoxide (diatomic, at room temperature) cluster tightly around 20.8. The theory is not approximately right; it is exactly right to within experimental precision.

Heat capacity at constant pressure

Now consider the right-hand cylinder in the figure. The piston is free to slide, and the weight on top keeps the pressure constant at whatever value the weight's mass and the piston's area deliver. Same gas, same temperature rise ΔT, different arrangement. The first law still holds, but this time W is not zero; the gas does expansion work against the weight and the atmosphere:

$$Q = \Delta U + W \quad \text{(constant P)}$$

The internal energy change ΔU depends only on the temperature change (for an ideal gas, U depends only on T, as Chapter 4 established), so ΔU is the same as in the constant-volume case for the same ΔT:

$$\Delta U = nC_v\Delta T$$

What about the work? Because the pressure is constant throughout, the work done by the gas is simply $W = P\Delta V$, from the previous section. For an ideal gas at constant pressure, the ideal-gas law PV = nRT, differentiated while holding P fixed, gives $P\Delta V = nR\Delta T$. So:

$$W = P\Delta V = nR\Delta T \quad \text{(constant P)}$$

Substituting both pieces into the first law:

$$Q = nC_v\Delta T + nR\Delta T = n(C_v + R)\Delta T$$

The **molar heat capacity at constant pressure** is defined, in parallel with the constant-volume version, as *the heat required per mole per degree at constant pressure:*

$$C_p = \frac{Q}{n\Delta T} \quad \text{at constant P}$$

Substituting the expression for Q:

$$C_p = \frac{n(C_v + R)\Delta T}{n\Delta T}$$

$$\mathbf{C_p = C_v + R \text{ or } C_p - C_v = R}$$

This is **Mayer's equation**, named for the German physician and physicist Julius Robert von Mayer who derived it in the 1840s, at about the same time Joule was pounding water with his paddle-wheel across the Channel. The equation is short, universal, and carries the whole payoff of the picture the section opened with. *The heat capacity at constant pressure exceeds the heat capacity at constant volume by exactly R, the universal gas constant*, no matter what gas is involved and no matter how many degrees of freedom its molecules have. The universal gas constant is precisely the expansion-work penalty per mole per kelvin for heating at constant pressure instead of constant volume.

Numerically, the predictions are:

$C_p = \frac{3}{2}R + R = \frac{5}{2}R = 20.8\ Jmol^{-1}\,K^{-1}$ for a monatomic gas and

$C_p = \frac{5}{2}R + R = \frac{7}{2}R = 29.1 Jmol^{-1}\,K^{-1}$ for a diatomic gas at room temperature.

Again the experimental values agree.

The adiabatic index

The ratio of the two molar heat capacities is denoted by the Greek letter gamma and called the **adiabatic index** of the gas:

$$\boldsymbol{\gamma = \frac{C_p}{C_v}}$$

For a monatomic ideal gas, $\gamma = (5/2)R / (3/2)R = 5/3 \approx 1.67$. For a diatomic ideal gas at room temperature, $\gamma = (7/2)R / (5/2)R = 7/5 = 1.40$. For polyatomic gases such as water vapour or carbon dioxide, γ is typically smaller still, around 1.33, because more degrees of freedom per molecule means a larger C_v and therefore a ratio closer to 1. *The value $\gamma = 1$ (which **no** real gas ever reaches) would mean the two heat capacities are equal, which by Mayer's equation would require $R = 0$,* which the universe has declined to arrange.

The ratio γ is not an idle number. It will reappear shortly as the exponent that distinguishes an adiabatic process from an isothermal one. An isothermal process obeys $PV = \text{constant}$, and an adiabatic process obeys $PV^{\gamma} = \text{constant}$, where the γ is exactly this ratio. The difference between slow compression and fast compression in the hand pump of the previous section is mediated by γ, and so γ deserves its careful introduction here, before the processes that use it.

The table below summarises the predictions of the theory for the three main categories of ideal gas. It is worth learning.

Gas type	f	C_v	C_p	γ	Examples
monatomic	3	(3/2)R	(5/2)R	5/3	He, Ar
diatomic	5	(5/2)R	(7/2)R	7/5	N_2, O_2
polyatomic	6 or more	3R	4R	4/3	H_2O, CO_2

Two rearrangements of Mayer's equation and the adiabatic index are worth deriving explicitly. Starting from:

$$\gamma = \frac{C_p}{C_v} \quad \text{and} \quad C_p = C_v + R,$$

substitute the second equation into the first to eliminate C_p:

$$\gamma = \frac{(C_v + R)}{C_v} = 1 + \frac{R}{C_v},$$

which rearranges to the first useful form:

$$\boldsymbol{C_v = \frac{R}{\gamma - 1}}$$

Multiplying both sides by γ and using $C_p = \gamma C_v$ gives the second:

$$\boldsymbol{C_p = \frac{\gamma R}{\gamma - 1}}$$

Both forms are handy when a problem gives γ and asks for either heat capacity, which is exactly the direction many examination problems take. They will also be the bridge by which the adiabatic equation $PV^{\gamma} = \text{constant}$ gets connected to the first law in the adiabatic section.

Specific heat capacities of gases

Everything derived so far in this section has been **molar**, meaning "per mole". Molar heat capacities are natural for gases because the equipartition theorem from Chapter 4 applies per molecule, and moles count molecules; so molar heat capacities come out in clean multiples of R/2. But many practical problems, give **specific** heat capacities instead, that is heat per kilogram per kelvin, with units of $Jkg^{-1}\,K^{-1}$. Chapter 2 used specific heat capacities throughout, for solids and liquids. This subsection bridges the two conventions for gases, and a student must be comfortable switching between them.

The specific heat capacity at constant volume, written c_v (lower-case c, to distinguish it from the molar C_v), is defined as *the heat needed per kilogram per degree temperature rise at constant volume*. The relationship to the molar capacity follows from noting that one mole of a gas has a mass equal to its molar mass M (in **kgmol**$^{-1}$). Thus:

$$m = \frac{n}{M}, \text{and } c_v = \frac{C_v}{M}$$

Analogously:

$$c_p = \frac{C_p}{M}$$

Dividing Mayer's equation by the molar mass on both sides:

$$\mathbf{c_p - c_v = \frac{R}{M}}$$

This is Mayer's equation in specific form. It says that *the specific heat capacity at constant pressure exceeds the specific heat capacity at constant volume by R/M,* which is the universal gas constant per unit mass.
A second form follows from the ideal-gas law.
From $PV = nRT$ with $n = \frac{m}{M}$, so $PV = \frac{mRT}{M}$, which rearranges to:

$$\frac{R}{M} = \frac{PV}{mT} = \frac{P}{\left(\frac{m}{V}\right)T} = \frac{P}{\rho T}$$

Substituting this back into the specific form of Mayer's equation gives:

$$\mathbf{c_p - c_v = \frac{P}{\rho T}}$$

This form is especially convenient when the problem gives the gas's density, pressure, and temperature, because it bypasses the molar mass entirely. Both forms say the same thing as Mayer's equation said in its molar form; all three are Mayer's equation, expressed in different units.
The adiabatic index γ is independent of which convention is used, because the molar mass M appears in both the numerator and the denominator of the ratio and cancels out:

$$\boldsymbol{\gamma = \frac{C_p}{C_v} = \frac{c_p}{c_v}}$$

This universality is why γ is the quantity most worth remembering: it has the same numerical value regardless of whether the data is expressed per mole or per kilogram.

A small rule for deciding which form of Mayer's equation to reach for. If the problem gives data in moles, use the molar form $C_p - C_v = R$. If the problem gives data in kilograms or via density, use whichever of the two specific forms is more direct: R/M when M is handed to you, and $P/(\rho T)$ when density, pressure, and temperature are given. Choosing the form that matches the given data is not stylistic; it avoids converting every number through molar mass, and it avoids the arithmetic error that conversions invite.

That is all the theory. Six worked examples follow.

BINDER Example 9

Two moles of helium are heated from 27°C to 127°C. Calculate the heat required (a) if the heating is done at constant volume, and (b) if the heating is done at constant pressure. Take $R = 8.314 Jmol^{-1} K^{-1}$.

Solution
The amount of substance is n = 2 mol. The temperature change is $\Delta T = 127°C - 27°C = 100K$ (the size of a degree is the same on the Celsius and the kelvin scales, so no conversion is needed). Helium is monatomic, so $C_v = (3/2)R$ and $C_p = (5/2)R$.
(a) Constant volume:

$$Q = nC_v\Delta T = n \times (3/2)R \times \Delta T = (2mol)(3/2)(8.314 Jmol^{-1}K^{-1})(100K) = 2494J$$

(b) Constant pressure:

$$Q = nC_p\Delta T = n \times (5/2)R \times \Delta T = (2mol)(5/2)(8.314 Jmol^{-1}K^{-1})(100K) = 4157J$$

Making Sense of the Answer: *Same gas, same amount, same temperature rise, two very different heat requirements. The constant-pressure case needs 1663J more heat than the constant-volume case, and Mayer's equation tells us where that extra heat went: into expansion work against the atmosphere, at a rate of R per mole per kelvin, totalling $nR\Delta T = (2mol)(8.314 Jmol^{-1}K^{-1})(100K) = 1663J$. The arithmetic closes the loop. The extra heat in the constant-pressure case is exactly the expansion work done.*

Think Like a Physicist: *A common examination trap is to quote a single "heat capacity of helium" and expect it to apply regardless of conditions. There is no such thing for a gas. Every calculation involving a gas must first decide whether the process is at constant volume, at constant pressure, or neither; and only after that decision can the correct heat capacity be selected. "Heat capacity at constant volume" and "heat capacity at constant pressure" are two different numbers, separated by exactly R, and choosing between them is a physics decision, not a formula look-up.*

BINDER Example 10

Compute the molar heat capacity at constant volume C_v, the molar heat capacity at constant pressure C_p, and the adiabatic index γ for (a) argon and (b) nitrogen. Express each value both in $Jmol^{-1}K^{-1}$ and as a multiple of R. Take $R = 8.314 Jmol^{-1} K^{-1}$.

Solution
(a) Argon is monatomic; f = 3, so the formulae give:

$$C_v = \left(\frac{3}{2}\right)\mathbf{R} = \frac{3}{2}\times(8.314\text{Jmol}^{-1}\text{K}^{-1}) = \mathbf{12.47Jmol^{-1}K^{-1}}$$

$$C_p = C_v + R = \left(\frac{3}{2}\right)R + R = \left(\frac{5}{2}\right)\mathbf{R} = \frac{5}{2}\times(8.314\text{Jmol}^{-1}\text{K}^{-1}) = \mathbf{20.79Jmol^{-1}K^{-1}}$$

$$\gamma = \frac{C_p}{C_v} = \frac{\left(\frac{5}{2}\right)R}{\left(\frac{3}{2}\right)R} = \frac{5}{3} = \mathbf{1.67}$$

(b) Nitrogen is diatomic; f = 5, so:

$$C_v = \left(\frac{5}{2}\right)\mathbf{R} = \frac{5}{2}\times(8.314\text{Jmol}^{-1}\text{K}^{-1}) = \mathbf{20.79Jmol^{-1}K^{-1}}$$

$$C_p = C_v + R = \left(\frac{5}{2}\right)R + R = \left(\frac{7}{2}\right)\mathbf{R} = \frac{7}{2}\times(8.314\text{Jmol}^{-1}\text{K}^{-1}) = \mathbf{29.10Jmol^{-1}K^{-1}}$$

$$\gamma = \frac{C_p}{C_v} = \frac{\left(\frac{7}{2}\right)R}{\left(\frac{5}{2}\right)R} = \frac{7}{5} = \mathbf{1.40}$$

Making Sense of the Answer: *A curious coincidence falls out: the constant-volume heat capacity of nitrogen, 20.79 J mol⁻¹ K⁻¹, is numerically equal to the constant-pressure heat capacity of argon. This is not accident. Going from argon (f = 3) to nitrogen (f = 5) adds two degrees of freedom worth of C_v, which is (2/2)R = R. And going from constant V to constant P also adds exactly R, by Mayer's equation. So (5/2)R for nitrogen at constant V equals (3/2)R + R = (5/2)R for argon at constant P. The gas constant R is everywhere in thermodynamics, and here it cancels two completely different physical effects against each other.*

Think Like a Physicist: *Memorising these six numbers is unnecessary if Mayer's equation and the (f/2)R rule are understood. Argon's C_v is (3/2)R; everything else in the list can be generated from that plus the two rules. C_p is C_v + R. Nitrogen's C_v is C_v of argon plus the extra (2/2)R from the two rotational modes. The adiabatic index is just the ratio. A student who has the two rules on the desk does not need to carry the table in memory; the table carries itself.*

BINDER Example 11

A gas has density 1.775 kgm⁻³ at 27°C and 10^5 Pa, and its specific heat capacity at constant pressure is c_p = 0.846 kJ kg⁻¹ K⁻¹. Find the specific heat capacity at constant volume c_v and the adiabatic index γ of the gas.

Solution

The data is given per kilogram (a density and a specific heat capacity), so we reach for the specific form of Mayer's equation rather than the molar form.

$$c_p - c_v = \frac{P}{\rho T} \text{ or } c_v = c_p - \frac{P}{\rho T}$$

Substituting:

$$c_v = 846\text{Jkg}^{-1}\text{K}^{-1} - \left(\frac{10^5\text{Pa}}{1.775\text{kgm}^{-3}\times 300\text{K}}\right) = 658\text{Jkg}^{-1}\text{K}^{-1} = 0.658\text{kJkg}^{-1}\text{K}^{-1}$$

The adiabatic index is the ratio:

$$\gamma = \frac{c_p}{c_v} = \frac{0.846\text{kJkg}^{-1}\text{K}^{-1}}{0.658\text{kJkg}^{-1}\text{K}^{-1}} = 1.29$$

Making Sense of the Answer: *A γ of about 1.3 is characteristic of a polyatomic gas with f = 6 active degrees of freedom (γ = 4/3 ≈ 1.33 exactly). Carbon dioxide, sulphur dioxide, and water vapour all sit in this neighbourhood. The small departure of 1.29 from the theoretical 1.33 is a real effect (no gas is ever perfectly ideal) and tells the physicist that the molecules are starting to interact with one another at this density and pressure. For examination purposes, 1.29 is close enough to 4/3, to classify the gas as polyatomic.*

Think Like a Physicist: *The first decision in every heat-capacity problem is not which equation to apply, but which form of the equation matches the given data. Molar data calls for the molar form. Kilogram data calls for the specific form. Density-pressure-temperature data calls for the P/(ρT) form. A student who learns to read the data before reaching for an equation will not waste time converting perfectly good numbers through a molar mass they did not need. The equations are all saying the same physics; the skill is matching form to data.*

REAL Example 12

A small *kioski* in Mbeya measures roughly 4m × 3m × 2.5m. On a cold morning the air inside is at 10°C; the owner wants to warm it to 22°C before opening for business, using a charcoal *jiko*. Take the density of air at 10°C and atmospheric pressure to be 1.25kgm⁻³, and its specific heat capacity at constant pressure to be c_p = 1005Jkg⁻¹K⁻¹. Calculate the heat the *jiko* must

deliver to the air. Why is the constant-pressure specific heat capacity the right one to use, even though the kioski is not a piston-and-cylinder arrangement?

Solution

The last question is where the physics is. A *kioski* is not a sealed box. Air leaks out under the door and around the windows as it warms, so the air inside never becomes pressurised above atmospheric. Throughout the heating, the pressure inside the *kioski* stays equal to the pressure outside, which is atmospheric. The air is being heated at constant pressure, and the correct heat capacity is therefore c_p. A piston is not needed; a leaky boundary does the same job.

With the physics settled, the arithmetic is short. The volume of the *kioski:*

$$V = 4\text{m} \times 3\text{m} \times 2.5\text{m} = 30\text{m}^3$$

The mass of air it contains:

$$m = \rho V = (1.25\text{kgm}^{-3})(30\text{m}^3) = 37.5\text{kg}$$

The temperature rise:

$$\Delta T = 22°\text{C} - 10°\text{C} = 12\text{K}$$

The heat required at constant pressure:

$$Q = mc_p\Delta T = (37.5\text{kg})(1005\text{Jkg}^{-1}\text{K}^{-1})(12\text{K}) = \mathbf{4.52 \times 10^5 J}\ (\approx 450\text{kJ}).$$

Making Sense of the Answer: *450kJ is the energy released by burning roughly 15 grams of charcoal (charcoal releases about 30MJ per kilogram when it burns). So the jiko needs to deliver 15 grams' worth of its fuel energy specifically to the air in the room. In practice the owner burns much more than 15 grams, because most of the jiko's heat escapes through its own flue, warms the pot sitting on top, or leaks into the walls and the floor rather than going into the air itself. The 450kJ is a lower bound, not an operational estimate. But it captures the fundamental physics that controls the morning's charcoal budget.*

Think Like a Physicist: *The kioski obeys constant pressure for the same reason an open cup of tea does and for the same reason the lungs do during a slow breath: any boundary that lets gas cross freely automatically imposes constant pressure on the gas inside, because excess pressure just pushes more gas out until the inside and outside agree. A piston is one way to keep pressure constant, but it is not the only way, or even the common way. Most constant-pressure processes in everyday life are constant-pressure because of leaky boundaries, not because of moving pistons. c_p, not c_v, is therefore the default heat capacity for heating a gas in an ordinary room or an ordinary pot. c_v applies only when the container is deliberately sealed, which, outside of a physics laboratory, is the exception rather than the rule.*

BINDER Example 13

The heat required to raise the temperature of 3.00mol of a polyatomic gas from 320K to 370K at constant pressure is 4.99kJ. Find (a) its molar heat capacity at constant pressure C_p, (b) its molar heat capacity at constant volume C_v, (c) its adiabatic index γ, and (d) the heat needed to raise 4.00mol of the same gas from 300K to 400K at constant volume. Take R = 8.314Jmol⁻¹K⁻¹.

Solution

(a) The molar heat capacity at constant pressure:

$$C_p = \frac{Q}{n\Delta T} = \frac{4990\text{J}}{3.00\text{mol} \times 50\text{K}} = 33.3\text{Jmol}^{-1}\text{K}^{-1}$$

This is very close to 4R = 33.26Jmol⁻¹K⁻¹, confirming the gas is polyatomic (f = 6).

(b) Apply Mayer's equation:

$$C_v = C_p - R = 33.3\text{Jmol}^{-1}\text{K}^{-1} - 8.314\text{Jmol}^{-1}\text{K}^{-1} = 25.0\text{Jmol}^{-1}\text{K}^{-1}$$

This is very close to 3R = 24.94Jmol⁻¹K⁻¹, again consistent with f = 6.

(c) The adiabatic index:

$$\gamma = \frac{C_p}{C_v} = \frac{33.3\text{Jmol}^{-1}\text{K}^{-1}}{25.0\text{Jmol}^{-1}\text{K}^{-1}} = 1.33$$

This matches the polyatomic prediction $\gamma = 4/3$ exactly.

(d) Now apply the calculated C_v to the new setup: n = 4.00mol, ΔT = 100K, constant volume. Since $\Delta V = 0$, all heat goes into internal energy:

$$Q = nC_v\Delta T = (4.00\text{mol})(25.0\text{Jmol}^{-1}\text{K}^{-1})(100\text{K}) = 1.00 \times 10^4\text{J} = 10.0\text{kJ}$$

Making Sense of the Answer: *Four numbers fall out consistently: $C_p = 4R$, $C_v = 3R$, $\gamma = 4/3$, and a final Q of 10.0kJ for the second setup. The first three match the polyatomic predictions to within the arithmetic's precision. The fourth applies the calculated C_v to a new question, which is what heat capacities are for*

Think Like a Physicist: *The flow of logic in this example is the flow of logic in every heat-capacity problem: measure whichever heat capacity the experimental setup gives you (here, C_p, because the setup was at constant pressure); use Mayer's equation to convert to the other (C_v); use the ratio to find γ; then apply either value to whatever new question arrives. Thermodynamics is economical with experimental data in a way that most of physics is not. Mayer's equation is the reason.*

HOT Example 14

A container holds 0.5 moles of nitrogen gas initially at 300K and 1.0×10^5Pa. The gas is heated so that its temperature rises by 200K. Calculate the heat supplied, the work done by the gas, and the change in internal energy, in each of the following three cases, and comment on the pattern:

(a) The heating is done at constant volume.
(b) The heating is done at constant pressure.
(c) The heating is done at constant temperature.

Solution

Nitrogen is diatomic gas. For a diatomic ideal gas at room temperature, $C_v = (5/2)R$ and $C_p = (7/2)R$. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$ and n = 0.5mol throughout.

(a) Constant volume, $\Delta T = 200\text{K}$. Since $\Delta V = 0$, no expansion work is done, and all the heat goes into internal energy:

$$W = 0$$
$$\Delta U = nC_v\Delta T = (0.5\text{mol})(5/2)(8.314\text{Jmol}^{-1}\text{K}^{-1})(200\text{K})\text{J} = 2079\text{J}$$
$$Q = \Delta U + W = 2079\text{J} + 0 = 2079\text{J}$$

(b) Constant pressure, $\Delta T = 200\text{K}$. Now the gas does expansion work, and the heat required is larger by a factor of $C_p / C_v = \gamma = 7/5 = 1.4$:

$$Q = nC_p\Delta T = (0.5\text{mol})(7/2)(8.314\text{Jmol}^{-1}\text{K}^{-1})(200\text{K})\text{J} = 2910\text{J}$$

The internal energy change is the same as in part (a), because ΔU depends only on ΔT for an ideal gas:

$$\Delta U = 2079\text{J}$$

The work done is the difference:

$$W = Q - \Delta U = 2910\text{J} - 2079\text{J} = 831\text{J}$$

As a consistency check, W should also equal $nR\Delta T$ for an ideal gas at constant pressure:

$$nR\Delta T = (0.5\text{mol})(8.314\text{Jmol}^{-1}\text{K}^{-1})(200\text{K})\text{J} = 831\text{J}\ \checkmark$$

(c) Constant temperature, $\Delta T = 0$. The internal energy of an ideal gas depends only on T, so if T does not change, neither does U:

$$\Delta U = 0$$

The first law then gives Q = W: whatever heat is supplied must emerge as work, and vice versa. Since $\Delta T = 0$, the formula $Q = nC_v\Delta T$ or $Q = nC_p\Delta T$ would return zero, but these formulae assume constant V or constant P, neither of which applies here. At constant temperature, both V and P must change together (along the hyperbola PV = nRT). So Q and W are not zero; they are equal to each other, and their value depends on how much the gas expanded or was compressed.

The calculation cannot be completed without additional information, and this is the point of including part (c): to show that the constant-volume and constant-pressure formulae are not general-purpose, and that the isothermal process needs its own treatment, coming in a later section.

Making Sense of the Answer: *The three cases form a pattern worth keeping in mind. At constant volume, all of Q goes into ΔU; W is zero. At constant pressure, some of Q goes into ΔU and the rest into W; the split is determined by Mayer's equation. At constant temperature, none of Q goes into ΔU; all of it goes into W. These are three of the four cornerstone processes of thermodynamics, and each one pins down a different one of the three thermodynamic variables (V, P, T) while the other two are allowed to adjust.*

Think Like a Physicist: *A gas can be heated along many different paths, and the heat required changes with the path even when the temperature change is the same. The gas does not care what label a physicist attaches to the process; it just responds to whichever of its state variables are being held fixed. The skill is to identify which variable is being held fixed in a given experimental setup, and then select the formula that corresponds to that constraint. At constant V, use C_v; at constant P, use C_p; at constant T, neither formula applies and the isothermal tools must be brought in.*

The two heat capacities now sit on the desk, connected by Mayer's equation, together with their ratio γ that will matter greatly for the adiabatic process. What remains is to put these tools to work on the four cornerstone processes of thermodynamics, taken one at a time. The first of these, the simplest by far, pins the volume down and lets everything else adjust. It is called the isochoric process, and it is the subject of the next section.

THE ISOCHORIC PROCESS (CONSTANT VOLUME)

Four thermodynamic processes now arrive, each with its own section. Each is defined by pinning one of the state variables down and letting the others adjust. The simplest of the four, and the first to meet, pins the volume. It is called the **isochoric process**, from the Greek **isos** (*equal*) and **choros** (*space*), and the picture that introduces it is probably already sitting on the reader's kitchen shelf.

Imagine a pressure cooker on a *jiko* in a kitchen. The lid is screwed firmly down; the safety valve is closed; the walls of the cooker are thick and stiff. Inside, a gas, perhaps the steam and air above a pot of beans being cooked. Below, the flame. As

the flame heats the gas inside, something a little surprising happens, or fails to happen. The walls of the cooker do not move. The volume of the gas inside does not change. Every molecule of the gas is now bouncing between walls that stay exactly where they were, at a distance it knows intimately. All that changes is how hard the molecules bounce. They bounce harder, because they are hotter. The pressure inside rises, but the volume does not.

This is the whole of the isochoric process, in one sentence. The gas inside is heated, but nothing moves. The picture on a PV diagram is therefore a single vertical line.

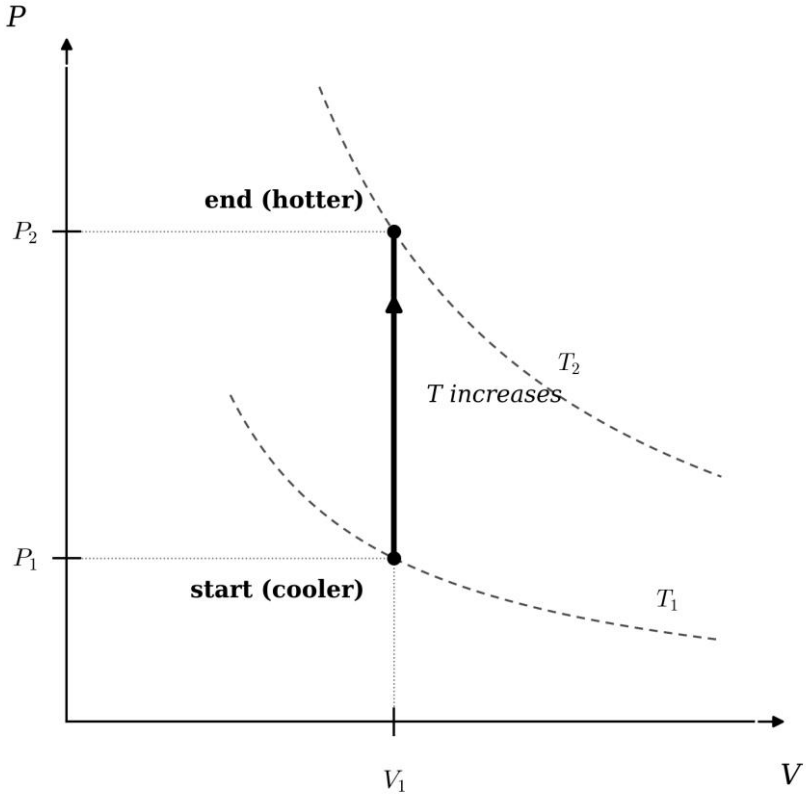


Figure: ***The isochoric process on a PV diagram.*** *The volume is held fixed at V_1 throughout the process, so the state of the gas moves only in the vertical direction, from the starting pressure P_1 (cooler) to the ending pressure P_2 (hotter). The two dashed curves are isotherms, representing the hyperbola PV = nRT at two different temperatures: T_1 through the lower endpoint and T_2 through the upper endpoint, with $T_2 > T_1$. The isochoric process carries the gas straight up from one isotherm to the other. No left or right motion is possible when the volume is fixed.*

The first law of thermodynamics, applied at constant volume

The defining feature of the isochoric process is $\Delta V = 0$. From the result of an earlier section, work done by a gas at constant volume is

$$W = \int P dV = 0 \quad (\text{since } dV = 0 \text{ at every point along the path}).$$

No volume change means no expansion against anything, and therefore no work. With W = 0, the first law simplifies dramatically:

$$Q = \Delta U + W = \Delta U + 0 = \Delta U$$

Every joule of heat supplied to the gas, in an isochoric process, goes into raising its internal energy. There is no other destination available to the heat, because the only two destinations are internal energy and work, and the work channel is shut. This is the isochoric process's chief feature, and the reason it is the simplest of the four processes: it has one channel where the others have two.

Using the Chapter 4 expression U = (f/2)nRT, the change in internal energy for a temperature change ΔT is:

$$\Delta U = \left(\frac{f}{2}\right) nR\Delta T = nC_v\Delta T$$

So the isochoric heat supply takes the convenient form:

$$\mathbf{Q = \Delta U = nC_v\Delta T} \quad \text{(isochoric)}$$

This is the first key result of the section, and it is what the student reaches for any time an isochoric calculation needs a heat.

How pressure and temperature relate during isochoric heating

A second result follows immediately from the ideal-gas law. For a fixed amount n of an ideal gas, PV = nRT at any moment during the process. If V is held fixed throughout, then n, R, and V are all constants, and the equation rearranges to:

$$P = \left(\frac{nR}{V}\right) \times T \text{ or } P = \text{constant} \times T,$$

which says P is proportional to T during the process. Equivalently, the ratio of P to T stays the same at every moment:

$$\frac{P}{T} = \text{constant} \quad \text{(isochoric)}$$

Alternatively, comparing the initial and final states of the process:

$$\frac{P_1}{T_1} = \frac{P_2}{T_2}$$

This is the isochoric form of the gas law, often called **Gay-Lussac's law** after the French chemist Joseph Louis Gay-Lussac, who published it about 220 years ago. It is a direct consequence of the ideal-gas law applied at fixed volume, and it is the relationship that any isochoric problem involving pressures and temperatures together will use.

The isochoric fingerprint

Three results together identify the isochoric process and make every calculation involving it a one-step exercise in substitution. They are worth memorising as a block:

$\mathbf{W = 0}$. No work, because no volume change.

$\mathbf{Q = \Delta U = nC_v\Delta T}$. All heat goes into internal energy.

$\mathbf{P/T = constant.}$ Pressure and absolute temperature rise together in fixed proportion.

Any problem involving a **rigid container**, a **sealed bottle**, a **pressure cooker** with the valve **shut**, a **gas cylinder** being heated in the sun, or any other situation in which a gas is heated or cooled without changing its volume can be read off these three results directly. The worked examples in this section walk through four such situations in sequence.

Nothing moves, everything rises, four examples now follow.

BINDER Example 15

A rigid container holds 0.32g of oxygen gas. The oxygen is heated from 25°C to 35°C. Given that the molar heat capacity at constant volume of oxygen is $C_v(O_2) = 20\text{Jmol}^{-1}\text{K}^{-1}$, and the molar mass of oxygen is $M(O_2) = 32\text{gmol}^{-1}$, find the heat supplied to the oxygen.

Solution

Convert the mass of oxygen to moles:

$$n = \frac{m}{M} = \frac{0.32\text{g}}{32\text{gmol}^{-1}} = 0.010\text{mol}$$

The temperature change is:

$$\Delta T = 35°\text{C} - 25°\text{C} = 10\text{K}$$

Since the container is rigid ($\Delta V = 0$), the isochoric heat formula applies:

$$Q = nC_v\Delta T = (0.010\text{mol})(20\text{Jmol}^{-1}\text{K}^{-1})(10\text{K}) = 2.0\text{J}$$

Making Sense of the Answer: *Two joules is a tiny amount of heat, which should not be a surprise: 0.32 grams is a tiny amount of oxygen (about the mass of a grain of rice), and ten degrees is a modest temperature rise.*

Think Like a Physicist: *The problem specified a "rigid container" deliberately. Without that phrase, the student would not know whether to reach for C_v or C_p, and the answer would be different by a factor of 7/5 for a diatomic gas like oxygen. Rigid means the volume is fixed, which means the process is isochoric, which means C_v is correct. Every word in a thermodynamics problem statement is doing work;* ***"rigid", "insulated", "open", "sealed",*** *and* ***"at constant pressure"*** *are the code words that identify which process is being described.*

BINDER Example 16

An insulated rigid container holds 2mol of helium gas. A heater inside the container supplies 60J of heat to the helium. Find the rise in temperature of the gas. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

Helium is monatomic, so $C_v = (3/2)R = 12.47\text{Jmol}^{-1}\text{K}^{-1}$. The container is rigid, so the process is isochoric, and C_v is the correct heat capacity. Apply the isochoric heat formula $Q = nC_v\Delta T$ and solve for ΔT:

$$\Delta T = \frac{Q}{nC_v} = \frac{60\text{J}}{2\text{mol} \times 12.47\text{Jmol}^{-1}\text{K}^{-1}} = 2.4\text{K}$$

The temperature rise is 2.4K.

Making Sense of the Answer: *The temperature rise is small, about two-and-a-half degrees, because sixty joules is not much heat and two moles of helium is a fair-sized sample. The phrase "insulated container" in the problem statement was not a tripwire; it simply confirmed that all the heat supplied by the heater went into the gas rather than leaking out through the walls. If the container had not been insulated, most of the sixty joules would have leaked away as heat to the surroundings, and the gas inside would barely have warmed at all. The insulation ensures the heat accounting closes cleanly: all 60J supplied, all 60J ending up in ΔU of the gas.*

Think Like a Physicist: *This is the inverse of the previous example: instead of finding Q given ΔT, find ΔT given Q. The same formula handles both directions. Whenever a heater is embedded in an insulated rigid container, the student should recognise it as the cleanest possible isochoric experiment: every joule delivered by the heater is a joule gained by the gas, and the temperature rise follows from C_v alone. Real laboratory measurements of C_v are made in exactly this kind of apparatus, called a bomb calorimeter, and the formula above is the equation used at the end of the measurement.*

REAL Example 17

A pressure cooker on a *jiko* contains air and steam at 20°C and 1.0atm when it is first sealed. The safety valve is set to open if the internal pressure reaches 2.5atm. Assuming the gas inside is an ideal gas, and taking the cooker's volume to be constant while the valve is still closed, find the temperature of the gas at the moment the safety valve opens.

Solution

While the valve is closed, the volume of the gas inside the cooker is fixed. The process is therefore isochoric, and the relation between pressure and absolute temperature is Gay-Lussac's law:

$$\frac{P_1}{T_1} = \frac{P_2}{T_2}$$

Temperatures must be in kelvin for this to work, because the ideal-gas law requires absolute temperature.

$$T_1 = 20°C + 273 = 293K$$

The pressures can be left in atmospheres, because they appear as a ratio and the units cancel. Rearrange Gay-Lussac's law for T_2:

$$T_2 = T_1 \times \frac{P_2}{P_1} = 293K \times \frac{2.5atm}{1.0atm} = 732.5K = 459.5°C$$

The temperature of the gas is $459.5°C$.

Making Sense of the Answer: *Four hundred and sixty degrees Celsius is hot. A cook heating a well-sealed pressure cooker on a jiko for long enough would indeed push the internal temperature into this range, which is precisely why pressure cookers have safety valves set to open well before such temperatures are reached. The student should be aware that real pressure cookers do not reach 459°C: long before that temperature is approached, either the food finishes cooking and the cook turns off the flame, or the safety valve opens and releases steam. The calculation shows what would happen in the worst case of a jammed valve, and is therefore a calculation safety engineers actually perform.*

Think Like a Physicist: *The isochoric step is the first half of the pressure cooker's story. The second half, which starts the instant the valve opens, is no longer isochoric: gas leaves through the valve, the volume of contained gas decreases, and the pressure stays pinned at 2.5atm by the valve mechanism. That second phase is isobaric, and it is the subject of the next section. Real devices rarely stay in a single thermodynamic process from start to finish; they typically walk through a sequence of processes, and the skilled student learns to recognise where one process ends and the next begins. The moment the valve opens in this problem is exactly such a transition.*

HOT Example 18

A rigid sealed pressure cooker contains 2.0mol of air at 20°C and 1.0atm. A heating element inside the cooker supplies energy at a steady rate of $500Js^{-1}$. The safety valve is set to open when the internal pressure reaches 2.5atm. Treating the air as a diatomic ideal gas, find:

(a) the time elapsed between the heater switching on and the moment the safety valve opens,
(b) the total heat supplied by the heater up to that moment, and
(c) the total change in internal energy of the gas up to that moment.
(d) Comment on the relationship between the answers to (b) and (c) in the light of the first law of thermodynamics.

Solution

Since air is diatomic, the molar heat capacity at constant volume is:

$$C_v = \left(\frac{5}{2}\right)R = \frac{5}{2} \times (8.314 Jmol^{-1}K^{-1}) = 20.79 Jmol^{-1}K^{-1}$$

Initial temperature in kelvin:

$$T_1 = 20°C + 273 = 293K.$$

Final temperature from Gay-Lussac's law:

$$T_2 = T_1 \times \frac{P_2}{P_1} = 293K \times \frac{2.5atm}{1.0atm} = 732.5K$$

Temperature change:

$$\Delta T = T_2 - T_1 = 732.5K - 293K = 439.5K$$

The heat required to produce this temperature change in the isochoric cooker:

$$Q = nC_v\Delta T = (2.0mol)(5/2)(8.314Jmol^{-1}K^{-1})(439.5K) = 18300J = 18.3kJ$$

(a) Time for the heater to deliver this much energy at $500Js^{-1}$:

$$t = \frac{Q}{rate} = \frac{18300J}{500Js^{-1}} = 37s$$

(b) The total heat supplied by the heater by that moment is Q itself:

$$Q = 18.3kJ$$

(c) The total change in internal energy is what the first law delivers after the work term has been evaluated. At constant volume, W = 0, so:

$$\Delta U = Q - W = 18.3kJ - 0 = 18.3kJ$$

(d) The two answers are numerically equal, both 18.3kJ. This is because the cooker is rigid, the gas cannot expand, so the work channel in the first law is closed: W = 0. With no work to account for, the first law $Q = \Delta U + W$ collapses to $Q = \Delta U$, and every joule of heat delivered by the heater ends up, with nothing lost to anything else, as an increase in the internal energy of the gas. This equality between Q and ΔU is the defining feature of every isochoric process.

Making Sense of the Answer: *The three numerical answers tell a consistent story. A 500-watt heater running for 37 seconds delivers 18.3kJ of heat, and that heat raises 2.0 moles of air from 293K to 732.5K in a rigid cooker whose volume holds fixed throughout. The pressure rises from 1.0 to 2.5 atmospheres in exact proportion to the temperature, as Gay-Lussac's law requires for an isochoric process. Every number checks against every other number; nothing is left over, and nothing is missing.*

Think Like a Physicist: *Rate problems appear throughout thermodynamics whenever the question asks "how long?" rather than "how much?" The recipe is always the same: compute the total energy required for the temperature change using the appropriate heat capacity, then divide by the supply rate to get the time. Heat capacity does the physics; division does the timing.*

That is the first of the four thermodynamic processes in hand. The volume was held fixed; the work channel was closed; Q and ΔU were equal. The next process in the sequence keeps the gas's behaviour a little livelier: the volume is allowed to change, but the pressure is now the variable that is held fixed. This is the isobaric process, and it is the subject of the next section.

THE ISOBARIC PROCESS (CONSTANT PRESSURE)

In the previous section the volume was held fixed. Now the pressure is held fixed instead. The second of the four thermodynamic processes is called the **isobaric process**, from the Greek *isos* (equal) and **baros** (*weight, pressure*), because the word originally referred to a gas heated with a fixed weight pressing on its piston. The weight imposes a fixed pressure on the gas below it; the pressure stays at that fixed value throughout the heating; the volume, now free to change, adjusts upwards to accommodate the expanding gas.

The proper imagination of this process is exactly the one the word describes: a vertical cylinder, capped above by a frictionless piston, loaded on top with a stone of exactly the right mass to keep the pressure inside the cylinder at atmospheric. A flame below heats the gas. The gas expands; the piston rises; the stone goes up with it. Nothing about the pressure changes: the stone is still the same stone, the atmosphere above is still the same atmosphere, and the gas inside is still at atmospheric pressure plus the small extra pressure the stone contributes. The volume changes; the temperature changes; the pressure does not.

An even more familiar picture you can form in your mind, and one already met earlier in this chapter, is an ordinary room. The Mbeya kioski of an earlier worked example was heated at constant pressure precisely because its boundary was leaky: air could exit under the door and around the windows as it warmed and expanded. The kioski never built up any excess pressure because the excess air simply left. A classroom, a kitchen, a bedroom on a cold morning in the southern highlands, all behave as leaky isobaric systems whenever the air inside is warmed. The piston-and-stone arrangement and the leaky room are two manifestations of the same physics: pressure pinned, volume free, and the first law of thermodynamics now managing two channels at once.

The picture on a PV diagram is a horizontal line, and the area beneath it is the work done.

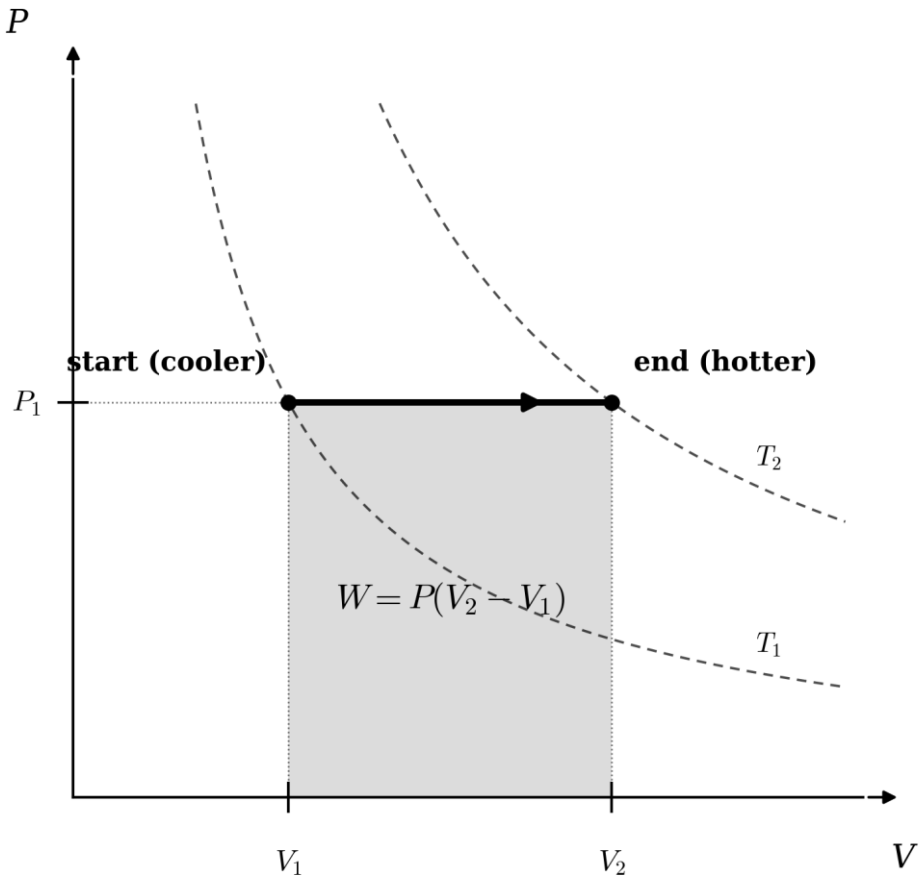


Figure: ***The isobaric process on a PV diagram.*** *The pressure is held fixed at P_1 throughout, so the state of the gas moves horizontally from the starting volume V_1 (cooler, on the lower isotherm T_1) to the final volume V_2 (hotter, on the upper isotherm T_2). The shaded rectangle below the horizontal line represents the work done by the gas during the expansion: $W = P(V_2 - V_1)$. The dashed curves are isotherms passing through the two endpoints, showing that as the gas moves to the right along the horizontal line, it is moving onto progressively hotter isotherms.*

The first law at constant pressure

The defining feature of the isobaric process is that the pressure is constant. From the result of an earlier section, work done by a gas at constant pressure is a single multiplication rather than an integral:

$$W = P(V_2 - V_1) = P\Delta V$$

This is the most direct isobaric result. But there is a more useful form, available because the gas is ideal. Start from the ideal-gas law at both states:

$$PV_1 = nRT_1, \quad PV_2 = nRT_2$$

Subtract the first equation from the second. Since P, n, R are all constant, they come out as common factors:

$$P(V_2 - V_1) = nR(T_2 - T_1) = nR\Delta T$$

So the isobaric work done has the clean alternative form:

$$\mathbf{W = P\Delta V = nR\Delta T} \quad \text{(isobaric)}$$

The right-hand form is often more convenient than the left, because it uses the temperature change directly and bypasses the volumes. Whichever form is faster in a given problem, it helps to know that both are available.

With the work term in hand, the first law gives the heat and the change in internal energy in parallel. The heat required at constant pressure is

$$Q = nC_p\Delta T \quad \text{(isobaric)},$$

which was already the defining relation for C_p in the molar heat capacities section. The change in internal energy, for any process of an ideal gas, depends only on the temperature change:

$$\Delta U = nC_v\Delta T \quad \text{(any process of an ideal gas)}.$$

This last formula is worth pausing on, because it looks surprising in an isobaric context. The gas **did** expand; the volume **did** change; and yet ΔU still uses C_v, the heat capacity at constant volume. The reason is that the internal energy of an ideal gas depends only on temperature, as the molecular derivation in Chapter 4 established. The C_v in $\Delta U = nC_v\Delta T$ is not saying

anything about constant volume; it is using the universal proportionality between U and T for an ideal gas. The proportionality constant happens to equal C_v, but that is a property of the gas, not a condition on the process.
With Q, ΔU, and W all expressed in terms of ΔT, the first law Q = ΔU + W is simply Mayer's equation in disguise:

$$\mathbf{nC_p\Delta T = nC_v\Delta T + nR\Delta T},$$

which divides through by nΔT to give $C_p - C_v = R$, the Mayer relation already familiar. This is a useful self-check: any isobaric calculation that tracks Q, ΔU, and W separately must satisfy this relation at the end.

How volume and temperature relate during isobaric heating

The ideal-gas law PV = nRT, applied at fixed pressure, has the same useful rearrangement that Gay-Lussac's law had at fixed volume. With P, n, R all constant, the equation reads:

$$V = \frac{nR}{P} \times T \text{ or } V = \text{constant} \times T,$$

which says V is proportional to T during the process. Equivalently:

$$\frac{V}{T} = \text{constant} \quad \text{(isobaric)},$$

or, comparing initial and final states,

$$\frac{V_1}{T_1} = \frac{V_2}{T_2}$$

This is known as **Charles's law**, after the French physicist Jacques Charles, who established it about 240 years ago through careful measurements on gases in open tubes. The law predates Gay-Lussac's by some twenty years, and historically it was Charles's observations that planted the idea of an absolute zero of temperature: extrapolating the V-versus-T line backwards, one finds that V would reach zero at around −273°C, which is no accident.

The isobaric fingerprint

Four results together identify the isobaric process. Any problem whose setup includes the words **"constant pressure", "isobaric", "open pot", "leaky room",** or a **piston loaded with a constant weight** can be solved by reaching for these directly:

$W = P\Delta V = nR\Delta T$. Work done by the gas, equal to the rectangular area under the horizontal process line.

$Q = nC\Delta T$. Heat supplied, using the constant-pressure molar heat capacity.

$\Delta U = nC_v\Delta T$. Change in internal energy, using the constant-volume molar heat capacity even though this is a constant-pressure process, because U depends only on T for an ideal gas.

$V/T = \text{constant}$. Volume and absolute temperature rise together in fixed proportion (Charles's law).

The piston rises, the stone rises, four examples now follow.

BINDER Example 19

2×10^{-2}kg of nitrogen gas is heated at constant pressure (1.0atm) so that its temperature rises by 45°C. Given $M(N_2) = 28\times10^{-3}$kgmol^{-1} and R = 8.314Jmol^{-1}K^{-1}, find the heat supplied.

Solution
The amount of substance in moles:

$$n = \frac{m}{M} = \frac{2\times10^{-2}\ \text{kg}}{28\times10^{-3}\ \text{kg mol}^{-1}} = 0.714\ \text{mol}.$$

Nitrogen is diatomic, so its molar heat capacity at constant pressure is:

$$C_p = \frac{7}{2}R = \frac{7}{2}\times 8.314\ \text{J mol}^{-1}\text{K}^{-1} = 29.10\ \text{J mol}^{-1}\text{K}^{-1}$$

The temperature rise is ΔT = 45K (the size of a kelvin equals the size of a degree Celsius). Applying the isobaric heat formula:

$$Q = nC_p\Delta T = (0.714\ \text{mol})(29.10\ \text{J mol}^{-1}\text{K}^{-1})(45\ \text{K}) = 935\ \text{J}$$

Making Sense of the Answer: *Nine hundred and thirty-five joules is a moderate amount of heat, equivalent to running a 100-watt bulb for about nine seconds, or burning roughly 0.03 grams of charcoal. Twenty grams of nitrogen gas warming by forty-five degrees is a small but not trivial heating job, and the 935J figure is appropriately sized. Had the process been carried out at constant volume instead, with $C_v = (5/2)R = 20.79$ J mol^{-1} K^{-1} in place of C_p, the heat needed would have been about 668J; some 267J less, exactly equal to nRΔT, which is the expansion work done by the gas in the constant-pressure case.*

Think Like a Physicist: *The phrase "heated at constant pressure" in an examination problem is a code word for isobaric, which in turn is a code word for "reach for C_p". This code-word recognition is what separates a confident thermodynamics student from a confused one. Whenever "constant pressure", "atmospheric pressure", "open container", or "unloaded piston" appears in a problem, the same mental switch must flip: isobaric, so use C_p.*

BINDER Example 20

A gas at 27°C and 10^5Pa has an initial volume of 0.04m³. It is heated at constant pressure until its volume becomes 0.05m³. The sample contains 48g of gas with molar mass 28gmol⁻¹. Find (a) the work done by the gas during the expansion, and (b) the final temperature of the gas.

Solution

(a) Work done by the gas at constant pressure:

$$W = P\Delta V = P(V_2 - V_1) = (10^5\ \text{Pa})(0.05\ \text{m}^3 - 0.04\ \text{m}^3) = 1000\ \text{J} = 1.0\ \text{kJ}$$

(b) For the final temperature, apply Charles's law:

$$\frac{V_1}{T_1} = \frac{V_2}{T_2}$$

Rearranging for T_2:

$$T_2 = T_1 \times \frac{V_2}{V_1}$$

With T_1 = 27°C + 273 = 300K:

$$T_2 = 300\ \text{K} \times \frac{0.05\ \text{m}^3}{0.04\ \text{m}^3} = 300\ \text{K} \times 1.25 = 375\ \text{K} = 102°\text{C}$$

Making Sense of the Answer: *A twenty-five-percent increase in volume produces a twenty-five-percent increase in absolute temperature, exactly as Charles's law requires for an isobaric process. The gas has gone from just below room temperature to just above the boiling point of water. Note that the problem provided more data than the two questions actually needed: the mass 48g and the molar mass 28gmol⁻¹ were not used in either part. This is common in physics problems and deliberate: real experimental setups give the physicist more measurements than any single calculation requires, and part of the skill is recognising which numbers matter for the question at hand. The moles n could have been computed from the given data and used in the alternative form W = nRΔT; both routes give the same answer.*

Think Like a Physicist: *The two formulae for isobaric work, W = PΔV and W = nRΔT, are both in hand from the derivation above. A careful student should notice that either form would work for part (a) of this problem. The first needs P and ΔV (both given directly). The second needs n (computed from m and M) and ΔT (which is what part (b) is asking for, so using this form would have created a circular dependency). The first form wins on efficiency. Whenever both forms of a formula are available, checking which data the problem hands you saves time: pick the form that takes raw inputs, not the form that needs a downstream result.*

REAL Example 21

A classroom with a radiator running for 10 minutes feels noticeably warmer at the end than at the start. The air in the room is at atmospheric pressure throughout, because the door, the windows, and the gaps around them are not airtight. Use the isobaric model to explain what happens to the volume of air as its temperature rises, where any "extra" air ends up, and why a well-sealed classroom would feel warmer than a leaky one for the same heat supplied by the radiator.

Solution

Take the air inside the classroom as the thermodynamic system. Because the door and windows are not airtight, any rise in internal pressure inside the room would immediately push air out through the gaps, so the pressure inside the room stays matched to atmospheric pressure throughout the heating. The process is therefore isobaric, and the relevant law is Charles's law:

$$\frac{V}{T} = \text{constant}$$

As the air's absolute temperature rises, its volume must rise in fixed proportion. But the classroom's physical dimensions, namely its walls, its floor, and its ceiling, do not move. The air cannot expand into extra space inside the room, because no extra space appears. Instead, air at the edges is pushed out through the leaks.

The air that remains inside is not the same air that was there at the start. It has been warmed by the radiator, some of its molecules have left through the leaks, and its density is correspondingly lower. Every cubic metre of classroom air now contains fewer molecules than it did at the start, because the same room volume now contains the same total number of molecules minus however many have leaked out. This is what warm air means: air at the same pressure but lower density than cold air.

A well-sealed classroom heats faster than a leaky one because all the radiator's heat stays in the room rather than walking out through the gaps. In a leaky room, each air molecule that escapes carries away a share of the internal energy the radiator has been supplying, so the energy that actually remains to warm the air inside is only a fraction of what the radiator delivered. In a sealed room, no molecules escape, and every joule the radiator delivers to the air stays as internal energy of that same air. For the same radiator running for the same time, the sealed room reaches a higher final temperature.

Making Sense of the Answer: *Three physical effects combine in the right way. Charles's law says volume rises with temperature at constant pressure. Fixed room dimensions say the volume of the enclosure cannot change. Leaky boundaries reconcile the two by letting the "extra" air exit through the gaps. The result is that a room at constant pressure and fixed geometry responds to heating by expelling air, not by expanding. This is not a defect of the isobaric model; it is the isobaric model applied to a real enclosed space whose boundaries are porous.*

Think Like a Physicist: *The principle that a leaky enclosure behaves isobarically is one of the most useful pieces of thermodynamic intuition a student can carry out of this chapter. It applies to a heated room, a heated lecture hall, a heated daladala interior with the windows open, a cooking pot with a loose lid, and the alveoli of the human lungs during a slow breath. A tightly sealed enclosure, by contrast, behaves isochorically: volume is fixed, no gas exits, and the pressure rises with temperature. The difference between the two extremes (and every case in between) is nothing more subtle than whether the boundary lets gas pass. Tight boundaries trap gas and raise pressure; leaky boundaries release gas and hold pressure.*

HOT Example 22

A cylinder contains 0.250mol of carbon dioxide at 17°C and 1.00atm. The gas is heated at constant pressure until its temperature reaches 127°C. Find (a) the work done by the gas, (b) the change in internal energy of the gas, (c) the heat supplied to the gas. (d) How much work would have been done if the constant pressure had been 0.50atm instead, with the same initial temperature and same final temperature? $R = 8.314 Jmol^{-1}K^{-1}$.

Solution

$$T_1 = 17°C + 273 = 290K, \qquad T_2 = 127°C + 273 = 400\ K, \Delta T = T_2 - T_1 = 400\ K - 290K = 110\ K$$

Since carbon dioxide is polyatomic, $\gamma = 1.33$

Calculating the two molar heat capacities:

$$C_v = \frac{R}{\gamma - 1} = \frac{8.314 Jmol^{-1}K^{-1}}{1.33 - 1} = 25.2 Jmol^{-1}K^{-1}$$

$$C_p = C_v + R = 25.2\ Jmol^{-1}\ K^{-1} + 8.314\ J\ mol^{-1}\ K^{-1} = 33.5\ J\ mol^{-1}\ K^{-1}$$

(a) Work done by the gas at constant pressure:

$$W = nR\Delta T = (0.250\ mol)(8.314\ J\ mol^{-1}\ K^{-1})(110\ K) = 229J$$

(b) Change in internal energy of the gas (depends only on ΔT for an ideal gas):

$$\Delta U = nC_v\Delta T = (0.250\ mol)(25.2\ J\ mol^{-1}\ K^{-1})(110\ K) = 693J$$

(c) Heat supplied, from the first law:

$$Q = \Delta U + W = 693J + 229J = 922J$$

As a consistency check, compute Q directly using C_p:

$$Q = nC_p\Delta T = (0.250\ mol)(33.5\ J\ mol^{-1}K^{-1})(110\ K) = 921J\ ✓$$

The two routes agree to within rounding, as they must for any isobaric process of an ideal gas.

(d) If the constant pressure had been 0.50atm instead, with the same T_1 and the same T_2, the work done would be:

$$W' = nR\Delta T = (0.250\ mol)(8.314\ J\ mol^{-1}\ K^{-1})(110\ K) = 229J$$

The work is the same. This is worth pausing on. At 0.50atm, the gas occupies twice the volume it occupied at 1.00atm (by Boyle's law applied to the initial state), and expanding at half the pressure through double the volume gives exactly the same PΔV. More formally, W = PΔV, but ΔV itself depends on P through the ideal-gas law: for the same ΔT, ΔV = nRΔT/P scales inversely with P, and the P in the numerator of PΔV cancels the P in the denominator of ΔV. The result is:

$$W = P \times \frac{nR\Delta T}{P} = nR\Delta T,$$

which contains no pressure at all. The isobaric work done on a given sample during a given temperature change is independent of the pressure at which that heating happens.

Making Sense of the Answer: *Three numerical answers and one surprise. The 229J of work, the 693J of ΔU, and the 922J of heat all fit together through the first law, and the two independent routes to Q agree. The surprise is that the work is the same regardless of the pressure, which contradicts an intuition students often bring in from the PV-diagram picture: half the pressure, the reader thinks, must give half the work. But the picture has to adjust along with the numbers. At 0.50atm the initial and final volumes are both twice what they were at 1.00atm, so the rectangle on the PV diagram is shorter in height and longer in width in exactly the right proportions to preserve its area. The rectangle's shape changes; its area does not. This is why reasoning from a formula sometimes delivers insights that a picture can miss.*

Think Like a Physicist: *The independence of isobaric work from pressure is an example of a more general principle: thermodynamic formulae are often more compact than their derivations, and they often reveal cancellations that are not obvious from the underlying physics. W = nRΔT carries no pressure on its face, even though W = PΔV does. The P drops out because ΔV itself depends on P. A student who internalises W = nRΔT as the "main" isobaric formula, with W = PΔV as its longer sibling, will move faster through calculations where only n and ΔT are at hand. More generally: when two forms of the same formula exist, the shorter one is usually the one that has absorbed cancellations the longer one still shows. Preferring the shorter form is a small but real time-saver on examination questions that must be answered against the clock.*

Two of the four thermodynamic processes now sit on the desk. The first held the volume fixed and shut the work channel. The second held the pressure fixed and opened both channels, with all the heat required at a rate set by C_p rather than C_v. The third process holds a different variable fixed, namely the temperature itself, and it unlocks a subtle point about how the first law manages heat and work when the internal energy of an ideal gas is not allowed to change. That process is called the isothermal process, and it is the subject of the next section.

THE ISOTHERMAL PROCESS (CONSTANT TEMPERATURE)

The previous section held the pressure fixed and let the volume breathe. This section holds the third and subtlest of the three primary state variables: the **temperature**. The process in which a gas's temperature remains unchanged throughout is called the **isothermal process**, from the Greek **isos** (*equal*) and **thermos** (*heat, temperature*). Holding the temperature fixed has a consequence the student may find surprising: the internal energy of the gas does not change either. U for an ideal gas depends only on T, and if T stays put, so does U. The first law then delivers a clean and slightly unusual result, Q = W, meaning every joule of heat flowing into the gas leaves again as work done by the gas on its surroundings. The gas is a pure energy conduit.

Imagine a gas sits inside a thin-walled cylinder. The cylinder is not an insulator; it is a conductor of heat. Surrounding the cylinder is a large reservoir of water, the kind one might find in a slow-moving river, or a bucket with a heating element carefully keeping its temperature fixed. A piston fits into the top of the cylinder. When the piston is pushed down very slowly, or pulled up very slowly, the gas inside does work against the piston, or has work done upon it. At every moment during this slow movement, whatever temperature change would tend to occur inside the gas is immediately undone by heat flowing into or out of the gas through the thin conducting wall. The reservoir is so large that no measurable heating or cooling of the reservoir itself occurs, no matter how long the experiment runs. The gas simply stays at the reservoir's temperature, and the heat bookkeeping records every joule.

Two conditions must hold for this description to apply. The cylinder must conduct heat freely, through thin walls, in constant contact with a large heat reservoir. And the process must be **slow** enough that heat has time to flow in or out at every instant along the way. A **quasi-static process,** the physicist calls it. A rapid compression in a thick-walled cylinder is the opposite of isothermal, and is the subject of the next section.

The picture on the PV diagram is a hyperbola.

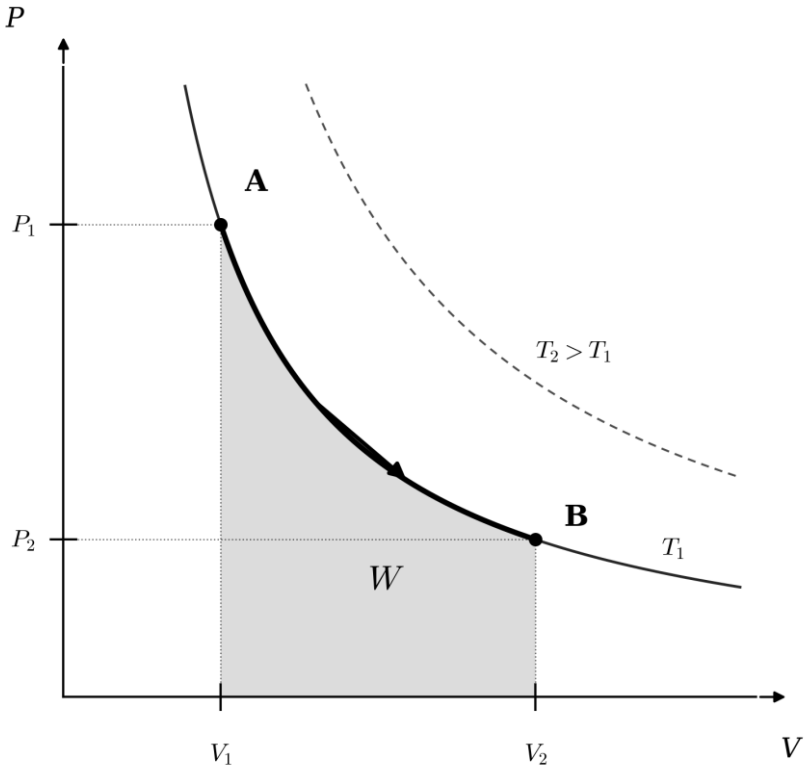


Figure: ***The isothermal process on a PV diagram.*** *The solid curve traces the hyperbola PV = constant at temperature T_1, carrying the gas from state A at (V_1, P_1) down-right to state B at (V_2, P_2) during an expansion. The shaded region beneath the curve between V_1 and V_2 represents the work done by the gas, W. A second isotherm, drawn dashed at a higher temperature $T_2 > T_1$, sits further from the origin: higher-T isotherms always lie further out than lower-T ones, because PV = nRT grows with T. The process shown is an expansion; the same curve, traced in the opposite direction from B to A, is an isothermal compression.*

The first law at constant temperature

For an ideal gas, the internal energy U depends only on the temperature T, as the molecular derivation in Chapter 4 established. If T is held fixed throughout a process, then

$$\Delta U = 0 \quad \text{(isothermal, ideal gas)}$$

The first law $Q = \Delta U + W$ simplifies at once:

$$\mathbf{Q = W} \quad \text{(isothermal, ideal gas)}$$

Every joule of heat that enters the gas leaves again as work done by the gas; every joule of work done on the gas must leave again as heat rejected to the surroundings. The gas's own energy store does not grow or shrink. This is a remarkable property: the gas acts as a perfect intermediary between heat and work, converting one to the other with no net change in its own state of excitement.

Boyle's law and the hyperbola

The second consequence of fixed temperature comes from the ideal-gas law. With T, n, and R all constant, PV = nRT becomes:

$$PV = \text{constant} \quad \text{(isothermal)},$$

or, comparing initial and final states,

$$P_1 V_1 = P_2 V_2$$

This is **Boyle's law**, named for the Anglo-Irish natural philosopher Robert Boyle, who published it about 360 years ago, when thermodynamics as a discipline was still 150 years away. Boyle measured the pressure and volume of a fixed amount of air at room temperature, confirmed that their product stayed constant as long as the temperature did, and handed physics a relationship that holds for every isothermal process of every ideal gas ever since.

On a PV diagram, PV = constant traces a hyperbola. Different temperatures give different hyperbolas, each called an **isotherm** for its particular temperature. Higher-T isotherms lie further from the origin, because for the same V, a higher T means a higher P, and for the same P, a higher T means a larger V. The dashed curve in the figure drawn at the higher temperature T_2, is one such isotherm; the solid curve, drawn at T_1, is another.

Work done during an isothermal process

The work done by the gas as its volume changes from V_1 to V_2 is the integral of PdV, just as in any process. What distinguishes the isothermal case is that P is not constant; it varies with V along the hyperbola, following P = nRT/V. Substituting into the work integral:

$$W = \int P\,dV = \int \frac{nRT}{V}\,dV$$

Since n, R, and T are all constants throughout the process, they come out of the integral:

$$\mathbf{W} = nRT \int_{V_1}^{V_2} \frac{dV}{V} = nRT[\ln V]_{V_1}^{V_2} = \mathbf{nRT \ln \frac{V_2}{V_1}}$$

Using Boyle's law $P_1V_1 = P_2V_2$ to exchange the volume ratio for a pressure ratio $\left(\frac{V_2}{V_1} = \frac{P_1}{P_2}\right)$ gives the equally useful alternative form:

$$\mathbf{W = nRT \ln \frac{V_2}{V_1} = nRT \ln \frac{P_1}{P_2}}$$

The two forms are mathematically identical. The first uses the volume ratio V_2/V_1; the second uses the pressure ratio P_1/P_2 (note the inversion; the subscripts swap, because when V goes up, P goes down). Either can be used, depending on which data the problem gives.

Since Q = W for an isothermal process of an ideal gas, the same formula delivers the heat:

$$\mathbf{Q = W = nRT \ln \frac{V_2}{V_1}}$$

What the sign of W tells us

Two cases arise, and the sign of W distinguishes them cleanly. If the gas expands, $V_2 > V_1$, so $V_2/V_1 > 1$, and $\ln(V_2/V_1) > 0$. Therefore, W > 0: positive work, done by the gas on its surroundings. The gas pushed back the piston; the surroundings absorbed the energy. And Q > 0 too: heat flowed in from the reservoir, exactly replacing what the gas spent as work.

If the gas is compressed, $V_2 < V_1$, so $V_2/V_1 < 1$, and $\ln(V_2/V_1) < 0$. Therefore, W < 0: negative work, as calculated. *What does a negative work mean physically?* It means the gas did not do positive work on its surroundings; rather, the surroundings did positive work on the gas; they pushed the piston down against the gas's resistance. The sign convention established in an earlier section defines W as work done by the gas on its surroundings. Negative W means the flow went the other way. And Q < 0 in lockstep: heat flowed out of the gas into the reservoir, carrying away exactly the work the surroundings had put in. The gas remained at its original temperature throughout, while acting as a pipe connecting the compressor's arm to the river's water.

Examiners like to ask, after a calculation of a negative W for an isothermal compression, "State the significance of the negative sign". The answer is not "the gas lost energy"; the gas did not lose energy, $\Delta U = 0$ throughout. The answer is **"the surroundings did work on the gas"** or equivalently, "the gas was compressed, not expanded". This distinction is worth owning.

The isothermal fingerprint

Four results identify every isothermal process of an ideal gas. You should be able to state them without looking at notes:

$PV = \text{constant}$. Boyle's law holds. The PV diagram traces a hyperbola.

$\Delta U = 0$. Internal energy depends only on T, and T does not change.

$Q = W$. Every joule of heat in leaves as work out; every joule of work in leaves as heat out.

$W = nRT\ln(V_2/V_1) = nRT\ln(P_1/P_2)$. Work done by the gas, equal to the shaded area under the hyperbola on the PV diagram.

The next **condition for applicability** is equally important: *the process must be slow, and the boundary must conduct heat freely*. Any process that violates either condition is not isothermal, and the formula above does not apply to it. The next section will formalize the opposite extreme, in which heat does not flow at all during the process.

The hyperbola is drawn, the logarithm is loaded, four examples now follow.

BINDER Example 23

One mole of an ideal gas at 320K is compressed isothermally from a volume of 8.0L to a volume of 4.0L. Find the work done by the gas and state the significance of the sign of your answer. Take R = 8.314Jmol^{-1}K^{-1}.

Solution

The process is isothermal, so:

$$W = nRT\ln\frac{V_2}{V_1} = (1.0\text{ mol})(8.314\text{ J mol}^{-1}\text{ K}^{-1})(320\text{ K})\ln\frac{4.0\text{ L}}{8.0\text{ L}} = -1844\text{ J} = -1.84\text{ kJ}$$

Significance of the negative sign. The work done by the gas is negative, which means the gas did **not** do positive work on its surroundings. Instead, the surroundings did 1.84kJ of positive work on the gas in order to compress it from 8.0L down to 4.0L. The gas was on the receiving end of the energy transfer, not the giving end.

Making Sense of the Answer: *The magnitude, about two kilojoules, is the energy cost of squeezing a mole of gas at room temperature to half its volume. In practice, this is the work a person must supply through a piston handle. Note that the gas's internal energy did not rise in response, because the process was isothermal; the heat reservoir absorbed the 1.84kJ of work as heat, and the temperature inside the cylinder never budged. If the reservoir had not been there, and the gas were compressed fast instead of slowly, the gas would have heated up and the calculation would have needed the adiabatic formulae of the next section.*

Think Like a Physicist: *Three things to own about the logarithm. First, ln(0.5) is negative, about −0.693; students sometimes forget this and return a positive answer. Second, the two conventions $\ln(V_2/V_1)$ for the isothermal work and the sign rule "positive W means work done by the gas" together guarantee that compression ($V_2 < V_1$) produces a negative W, automatically, without any extra sign juggling. Third, the formula's dependence on n and T is linear: compressing two moles gives twice the work; doing it at twice the temperature gives twice the work; halving the volume always gives $-nRT \times \ln(2) = -0.693nRT$ of work, whatever n and T happen to be. These scalings are worth internalizing.*

BINDER Example 24

Sixteen moles of an ideal gas at 320K are compressed isothermally from a volume of 18.0L to a volume of 4.0L. Find the work done by the gas and explain the sign. Take R = 8.314Jmol^{-1}K^{-1}.

Solution

The process is isothermal. Again we apply $W = nRT\ln(V_2/V_1)$:

$$W = (16\text{ mol})(8.314\text{ J mol}^{-1}\text{ K}^{-1})(320\text{ K})\ln\frac{4.0\text{ L}}{18.0\text{ L}} = -64020\text{ J} = -64.0\text{ kJ}$$

The sign is negative. The surroundings did 64kJ of work on the gas in order to compress it from 18L down to 4L. At the same time, 64kJ of heat flowed out of the gas into the surrounding reservoir, because Q = W for an isothermal process. The gas was a short-term stopover for the energy; the surroundings' work handle was pushing energy in at one side and the reservoir was receiving it as heat on the other.

Making Sense of the Answer: *Sixteen times the moles of the previous example, plus a larger volume ratio (4.5 rather than 2), gives a much larger work: 64kJ instead of 1.84kJ. Both scalings check out. More gas means more molecules to push against each other's pressure during compression; a larger compression ratio means the work integral extends over a larger*

range. The linearity in n is exact; the dependence on V_2/V_1 is logarithmic, so squeezing a gas down to a quarter of its volume does not cost four times.

Think Like a Physicist: *A quick way to spot a sign error in an isothermal problem: if the gas was compressed, W must be negative, and Q must be negative too (heat leaving the gas). If the gas expanded, both must be positive (heat entering, work done by the gas on surroundings). Any answer where the process direction and the sign of W disagree is a calculation error, either in the ratio V_2/V_1 or in the logarithm's sign. Check before moving on.*

REAL Example 25

A car tyre suddenly bursts on the Morogoro road. Kipute, who is standing nearby, notices that the air rushing out of the hole feels noticeably cold, not warm. Puzzled, she asks Mr. Akilikubwa about it later in the laboratory.

Kipute: *"Mr. Akilikubwa, when the tyre burst, the escaping air was cold. I thought compressed air should be warmer than the surroundings, not colder. Why did it feel cold?"*

Mr. Akilikubwa: *"A good question. The air inside the tyre was at about 2 atmospheres of pressure and at roughly the temperature of the road surface, which on a hot afternoon in Morogoro might be 40°C. When the tyre burst, the air expanded very quickly to atmospheric pressure. That expansion did work against the atmosphere outside, pushing air out of the way to make room for itself. **Where did that work-energy come from**? From inside the gas itself, because the burst was so fast that no heat from the tyre walls or the surrounding air had time to flow into the escaping gas during the expansion. With no heat coming in, and work going out, the gas's internal energy fell; and for an ideal gas, falling internal energy means falling temperature. The air cooled as it escaped, by many tens of degrees. That cold rush is what you felt."*

Kipute: *"So the compressed air wasn't warmer, it was cooler by the time it reached me."*

Mr. Akilikubwa: *"Yes. Two temperatures at play: the tyre-wall temperature, which never changed in the brief instant, and the air-inside-the-burst temperature, which dropped sharply during the expansion because there was no time for heat to flow. That quick-expansion process, where the gas cannot exchange heat with its surroundings, has a name. It is called an **adiabatic** process, and it is the subject of our next lesson. The first law still applies: $Q = \Delta U + W$, but with $Q = 0$ because no heat flowed."*

Kipute: *"Interesting. Now, what if the tyre had been deflated slowly? Say, through a pinhole so small that the air took hours to escape. Would it still have cooled?"*

Mr. Akilikubwa: *"Ah, this is where the isothermal picture we have been studying comes in. If the deflation is slow enough; very slow, hours-slow, then during the expansion, heat has time to flow from the tyre wall, from the road, from the surrounding air, into the escaping gas at every moment. The gas draws in just enough heat to replace the energy it is spending as work against the atmosphere. Its temperature therefore stays roughly equal to the temperature of its surroundings throughout the expansion. No dramatic cooling; no dramatic warming. The gas emerges at the ambient temperature, having quietly traded heat from the road for work done against the air. That is the isothermal case, and the formula we have just derived, $W=nRTln(V_2/V_1)$, computes the work exactly."*

Kipute: *"So fast expansion cools the gas; slow expansion keeps the temperature constant. The speed of the process picks which one."*

Mr. Akilikubwa: *"Precisely. And the physics that separates them is not the gas itself; it is the time available for heat to flow through the walls of whatever contains the gas. Thick walls, fast process, little time: adiabatic. Thin walls, slow process, plenty of time: isothermal. Every gas expansion in the world sits somewhere between these two idealizations, and the physicist's first job is to decide which one it is closer to."*

Making Sense of the Answer: *The same physical act: air escaping a tyre; has two completely different outcomes depending on the speed. Fast: adiabatic, gas cools sharply. Slow: isothermal, gas stays at ambient temperature. The two limits are the two thermodynamic idealizations between which real processes live. The burst tyre was adiabatic (bang, no time for anything); a slow pinhole leak would have been isothermal (hours, all the time in the world).*

Think Like a Physicist: *Real-world gas expansions and compressions are never exactly isothermal or exactly adiabatic, but almost all of them are close to one or the other. The classification is fast: if the time for the process is much shorter than the time for heat to flow through the boundary, the process is effectively adiabatic. If the time for the process is much longer, it is effectively isothermal. The ratio of the process time to the heat-flow time determines which idealization applies. Students who learn to make this call quickly are students who get thermodynamics right on the first try, because they have already chosen the correct formula before touching the calculator.*

HOT Example 26

Two moles of an ideal gas are compressed isothermally from 10L to 5L. Perform the calculation at two different temperatures: (a) T = 300K, and (b) T = 600K. For each case, find the work done on the gas and the heat released by the gas. (c) Comment

on the ratio of the two works and the physical reason behind the scaling. (d) For the hotter case, by what factor does the ratio V_2/V_1 need to change for the work to match the colder case's answer? Take R = 8.314Jmol^{-1}K^{-1}.

Solution

(a) At T = 300K:

$$W_a = nRT \ln \frac{V_2}{V_1} = (2.0\ \text{mol})(8.314\ \text{J mol}^{-1}\ \text{K}^{-1})(300\ \text{K}) \times \ln \frac{5\text{L}}{10\text{L}} = -3458\ \text{J} \approx -3.46\ \text{kJ}$$

The work is −3.46kJ, so the surroundings do 3.46kJ of positive work on the gas. Since Q = W for an isothermal process:

$$Q_a = -3.46\ \text{kJ},$$

meaning 3.46kJ of heat leaves the gas into the surroundings during the compression.

(b) At T = 600K, with the same n and same volume ratio:

$$W_b = (2.0\ \text{mol})(8.314\ \text{J mol}^{-1}\ \text{K}^{-1})(600\ \text{K}) \times \ln \frac{5\text{L}}{10\text{L}} = -6916\ \text{J} \approx -6.92\text{kJ}$$

And the heat released:

$$Q_b = -6.92\ \text{kJ}$$

(c) The ratio of the two works:

$$\frac{W_b}{W_a} = \frac{-6.92\ \text{kJ}}{-3.46\ \text{kJ}} = 2.00.$$

Comment: The work at 600K is exactly twice the work at 300K.

Physical reason: The isothermal work formula W = nRT ln(V_2/V_1) is linear in T. Doubling T doubles W for fixed n and fixed volume ratio. Compressing hot gas costs twice as much work as compressing the same amount of the same cold gas through the same volume ratio, because at the higher temperature the molecules are banging harder (with higher KE) against the piston and the piston must push harder to displace them.

(d) For the hotter case at 600K, we ask: *what **volume ratio** would make the work match the colder case at 300K, namely −3.46kJ?* Set up the equation:

$$(2.0\ \text{mol})(8.314\ \text{J mol}^{-1}\ \text{K}^{-1})(600\ \text{K}) \times \ln(\text{ratio}) = -3458\ \text{J},$$

which rearranges to:

$$\ln(\text{ratio}) = \frac{-3458\ \text{J}}{9977\ \text{J}} = -0.3466,$$

$$\text{ratio} = e^{-0.3466}$$

$$\frac{V_2}{V_1} = e^{-0.3466} = 0.707$$

So at 600K, the gas need only be compressed to about 71% of its original volume (not to 50%, as in the colder case) to lose the same 3.46kJ of work.

Making Sense of the Answer: *Parts (a), (b), and (c) together show the linear-in-T scaling of isothermal work in action: doubling the temperature doubles the work at the same volume ratio. Part (d) inverts the question: given a target work, how much compression does a hotter gas need? Answer: less.*

Think Like a Physicist: *The isothermal formula is linear in n, linear in T, and logarithmic in the volume ratio. Three different scalings, three different meanings. A doubling of n doubles the work (more molecules, each contributing the same per-molecule expansion energy). A doubling of T doubles the work (same molecules, each pushing harder). A doubling of the volume ratio adds ln(2) × nRT to the work (logarithmic, so ratios add in log-space rather than multiplying). Students who carry these three scalings in their head can estimate any isothermal problem's answer to within a factor of two before they reach for a calculator, which is exactly the speed an examination demands.*

Three of the four thermodynamic processes now sit on the desk. The volume was held fixed; the pressure was held fixed; the temperature was held fixed. The fourth and last process holds something less obvious: the heat exchange itself is held at zero. In such a process the gas is thermally sealed from its surroundings, and any change of volume must be paid for from the gas's own internal energy, with consequences for the temperature that the isothermal case kept carefully disguised. That process is the adiabatic process, and it is the subject of the next section, where the dialogue of REAL Example 25 will receive its formal treatment.

THE ADIABATIC PROCESS (NO HEAT EXCHANGE)

The previous section held the temperature constant by allowing heat to flow freely through a thin conducting wall, in a process slow enough for that flow to keep up at every instant. The fourth and last of the basic thermodynamic processes goes to the opposite extreme. Instead of holding T constant, it holds the heat exchange Q at zero. Either the gas is wrapped in walls so **insulating** that no heat can flow through them in the available time, **or** the process is so **fast** that heat has no time to flow even through ordinary walls. Either way, Q = 0, and the gas exchanges energy with the outside world purely as mechanical

work. The process is called the **adiabatic process**, from the Greek **a** (not) and **dia-bainein** (*to pass through*), in honour of the heat that does not pass through.

Imagine a cylinder with thermos-flask walls and a piston at the top. The walls keep heat out as well as a vacuum-jacketed flask keeps a hot drink hot, which is to say, very nearly perfectly. When the gas inside is compressed by pushing the piston down, the work done on the gas has nowhere to go except into raising its internal energy, which raises its temperature. When the gas is allowed to expand by letting the piston rise, the work the gas does on the piston comes out of its internal energy, which falls, taking the temperature with it. The bicycle pump that warmed up under fast pumping in Chapter 4, and the burst tyre whose escaping air felt cold in the dialogue of an earlier section, are both examples of this process happening in the wild. Fast processes do not give heat time to leak through walls; they are adiabatic by virtue of speed alone, even when the walls are not particularly insulating.

On a PV diagram, the adiabat is a curve that resembles the isotherm but is steeper. The figure below shows the two together, drawn from a common starting point so the comparison is direct.

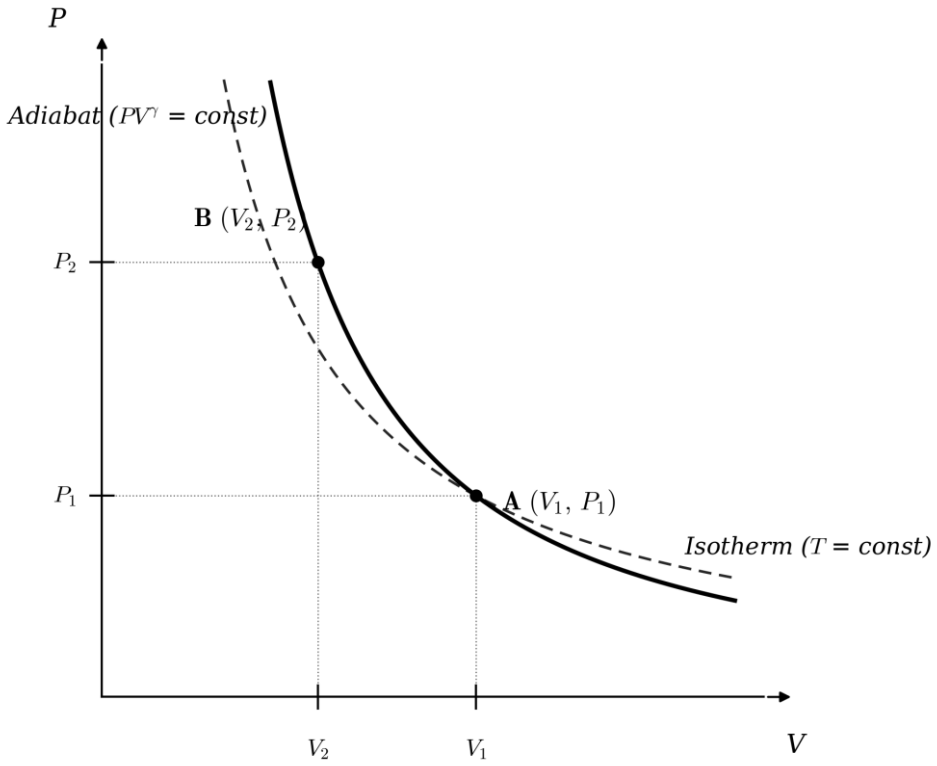


Figure: ***An adiabatic process and an isothermal process drawn from the same initial point A.*** *The adiabat (solid) follows PV^{γ} = constant; the isotherm (dashed) follows PV = constant. The adiabat is everywhere steeper than the isotherm: when the gas is compressed adiabatically from A to B, both V and T decrease less than they would along the isotherm, so the pressure rises faster per unit of volume change. The physical reason is the first law: an isothermal compression sheds heat to the surroundings as it goes, keeping T fixed; an adiabatic compression keeps that heat in, so T rises during the process and the rising T contributes its own pressure rise on top of the geometric one from compression. The adiabat is therefore steeper than the isotherm by exactly the factor γ, the adiabatic index.*

The first law with no heat exchange

With Q = 0, the first law $Q = \Delta U + W$ rearranges immediately:

$$\Delta U = -W \quad \text{(adiabatic)}$$

This is the central physics of the adiabatic process, and it deserves a moment to absorb. In every other process so far, work was a side effect of moving energy between the heat reservoir and the gas. In an adiabatic process, work and internal energy are directly tied: every joule of work the gas does on its surroundings is a joule taken out of its internal energy, and every joule of work done on the gas by its surroundings is a joule added to its internal energy. There is no third party, no heat reservoir, no thermometer-fixing river. The gas's energy store is the only buffer.

Since the internal energy of an ideal gas depends only on temperature, $\Delta U = nC_v\Delta T$, the equation rearranges to a temperature statement:

$$\mathbf{nC_v\Delta T = -W},$$

which says: *when work is done by the gas (W > 0), ΔT < 0 and the gas cools; when work is done on the gas (W < 0), ΔT > 0 and the gas warms.* The cold air rushing from the burst tyre, the warm bicycle-pump barrel, and the warm hands rubbed in fast friction all share this single mechanism.

Derivation of the adiabatic equation

The first law in differential form, with Q = 0, gives a relation between an infinitesimal volume change and an infinitesimal temperature change. The work done by the gas during a small expansion dV is PdV, and the change in internal energy during a small temperature change dT is $nC_v dT$. Setting $\Delta U = -W$ in their differential forms:

$$nC_v dT = -PdV \quad (1)$$

This is one equation in two infinitesimal unknowns dT and dV. A second equation comes from the ideal-gas law PV = nRT. Differentiating both sides (the product rule on the left, only T varying on the right with n and R constant):

$$PdV + VdP = nRdT \quad (2)$$

Now eliminate dT between (1) and (2). From (1):

$$dT = -\frac{PdV}{nC_v}$$

Substituting into (2):

$$PdV + VdP = nR \times \left(-\frac{PdV}{nC_v}\right) = -\frac{R}{C_v}PdV$$

Move the right-hand-side term to the left and combine the PdV terms:

$$PdV + \frac{R}{C_v}PdV + V\,dP = 0,$$

$$\left(1 + \frac{R}{C_v}\right)PdV + VdP = 0$$

The factor in parentheses simplifies neatly. Using Mayer's equation:

$C_p - C_v = R$, the ratio $\frac{R}{C_v} = \frac{C_p - C_v}{C_v} = \frac{C_p}{C_v} - 1 = \gamma - 1$, so $1 + \frac{R}{C_v} = 1 + \gamma - 1 = \gamma$. The equation becomes:

$$\gamma PdV + VdP = 0 \quad (3)$$

Pause for a moment to look at what has just happened. The temperature has been eliminated. Equation (3) involves only P, V, and the constant γ; no T at all. This is genuinely surprising. We started from a first-law equation that featured the temperature change dT, and we ended with a relationship purely between P, V, and dP, dV. The temperature has been hidden inside γ, where it is no longer a variable but a fixed property of the gas. The relationship that survives is between pressure and volume only. Students often find this elimination step disorienting on first reading; it is worth re-reading until the surprise has worn off, because the next step turns equation (3) into the famous adiabatic equation commonly examined in advanced thermodynamics problems.

Divide equation (3) through by PV (this assumes the gas has non-zero pressure and volume, which is always true for a real gas in a real situation):

$$\gamma\frac{dV}{V} + \frac{dP}{P} = 0$$

Each term is now exact: dV/V is the differential of ln V, and dP/P is the differential of ln P. Integrating both sides:

$$\int \gamma\frac{dV}{V} + \int \frac{dP}{P} = \int 0$$

$$\gamma \ln V + \ln P = \text{constant or } \ln V^{\gamma} + \ln P = \text{constant}$$

Combining logarithms:

$$\ln(PV^{\gamma}) = \text{constant},$$

which means the argument of the logarithm is itself a constant:

$$\mathbf{PV^{\gamma} = constant} \quad \text{(adiabatic)}$$

This is the **adiabatic equation**, and it is the central result of this section. *For any **quasi-static adiabatic process** of an ideal gas, the product P times V to the power γ stays constant from the start of the process to the end.* Or, comparing initial and final states using subscripts on each variable:

$$\mathbf{P_1 V_1^{\gamma} = P_2 V_2^{\gamma}} \quad (4)$$

Two companion forms

The adiabatic equation (4) connects pressure and volume. Two further relations connect temperature with each of them, derived by substituting the ideal-gas law into (4).

First, eliminate P using P = nRT/V. Equation (4) becomes:

$$\frac{nRT_1}{V_1} V_1^{\gamma} = \frac{nRT_2}{V_2} V_2^{\gamma}$$

The nR cancels, and the V^{γ}/V combines into V to the power $(\gamma - 1)$:

$$\mathbf{T_1\, V_1^{\gamma-1} = T_2\, V_2^{\gamma-1}}$$

which is often written compactly as:

$$\mathbf{TV^{\gamma-1} = constant} \quad \text{(adiabatic)}$$

This is the form of choice when V is what is changing during the process and T is what is being asked for.
Second, eliminate V using V = nRT/P. Equation (4) becomes:

$$P_1\left(\frac{nRT_1}{P_1}\right)^{\gamma} = P_2\left(\frac{nRT_2}{P_2}\right)^{\gamma}$$

The constants nR cancel between the two states, and after simplification:

$$\frac{T_1^{\gamma}}{P_1^{\gamma-1}} = \frac{T_2^{\gamma}}{P_2^{\gamma-1}}$$

This rearranges to a particularly useful form for problems where pressure changes adiabatically:

$$\left(\frac{T_2}{T_1}\right)^{\gamma} = \left(\frac{P_2}{P_1}\right)^{\gamma-1}$$

$$\mathbf{\frac{T_2}{T_1} = \left(\frac{P_2}{P_1}\right)^{\frac{(\gamma-1)}{\gamma}}}$$

Writing it in this fraction-and-exponent form lets the student read off the temperature ratio directly from the pressure ratio. This is the form the worked examples below will reach for most often.

Three forms of the adiabatic work, all equivalent

The work done by the gas during an adiabatic process can be written in three equivalent forms. Each is convenient in a different context, and each can be derived from a different starting point. The derivations are short but worth seeing in full.

Form 1: from the work integral

Start from the general definition of work done by a gas during a quasi-static process:

$$W = \int_{V_1}^{V_2} P\,dV$$

During an adiabatic process, the pressure depends on volume through PV^{γ} = constant. Setting that constant equal to its value at the initial state, $P_1 V_1$ raised to γ:

$$P\,V^{\gamma} = P_1\,V_1^{\gamma},$$

which rearranges to express the pressure during the process in terms of V:

$$P = \frac{P_1\,V_1^{\gamma}}{V^{\gamma}} = P_1\,V_1^{\gamma}\,V^{-\gamma}$$

The factor $P_1\,V_1^{\gamma}$ with subscript 1 is constant during the process; only V varies. So:

$$W = \int_{V_1}^{V_2} P\,dV = P_1\,V_1^{\gamma}\int_{V_1}^{V_2} V^{-\gamma}\,dV = P_1\,V_1^{\gamma}\left[\frac{V^{1-\gamma}}{1-\gamma}\right]_{V_1}^{V_2} = P_1\,V_1^{\gamma} \times \frac{\left(V_2^{1-\gamma} - V_1^{1-\gamma}\right)}{1-\gamma}$$

Multiply through and simplify each term:

$$W = \frac{P_1\,V_1^{\gamma}V_2^{1-\gamma} - P_1\,V_1^{\gamma}V_1^{1-\gamma}}{1-\gamma} = \frac{P_1\,V_1^{\gamma}V_2^{1-\gamma} - P_1\,V_1}{1-\gamma}$$

The term $P_1\,V_1^{\gamma}V_2^{1-\gamma}$ can be simplified further by using adiabatic relation: $P_1\,V_1^{\gamma} = P_2\,V_2^{\gamma}$

$$W = \frac{P_2\,V_2^{\gamma}V_2^{1-\gamma} - P_1\,V_1}{1-\gamma} = \frac{P_2\,V_2 - P_1V_1}{1-\gamma}$$

Multiplying numerator and denominator by −1 to put the formula into its conventional form (positive denominator):

$$\mathbf{W = \frac{P_1\,V_1 - P_2\,V_2}{\gamma - 1}} \quad \text{(form 1)}$$

This is the adiabatic work in pressure-volume form. It can be applied directly whenever the initial and final pressures and volumes are known.

Form 2: from form 1 via the ideal-gas law

The ideal-gas law applied at each end of the process gives:

$$P_1\,V_1 = nRT_1, \quad P_2\,V_2 = nRT_2$$

Substitute these into form 1:

$$\mathbf{W = \frac{nRT_1 - nRT_2}{\gamma - 1} = \frac{nR(T_1 - T_2)}{\gamma - 1}} \quad \text{(form 2)}$$

This is the temperature form. It is the form to use when initial and final temperatures and the moles are at hand, but volumes or pressures are not.

Form 3: from the first law directly

The first law of thermodynamics with Q = 0 gives ΔU = −W, which rearranges to W = −ΔU. The internal energy of an ideal gas depends only on temperature:

$$\Delta U = nC_v(T_2 - T_1)$$

So:

$$\mathbf{W = -nC_v(T_2 - T_1) = nC_v(T_1 - T_2)} \quad \text{(form 3)}$$

This is the internal-energy form. It is the fastest way to compute W when ΔU has already been computed for some other reason, and it provides an independent route to the answer that can verify forms 1 or 2.

How the three forms reconcile

Forms 1 and 2 differ only by the substitution PV = nRT, applied at each state. They are mathematically identical for any ideal gas.

Forms 2 and 3 differ only by Mayer's equation. From the rearrangements derived in the molar heat capacities section, $C_v = R/(\gamma-1)$. Substituting into form 3:

$$W = nC_v(T_1 - T_2) = n \times \frac{R}{\gamma - 1} \times (T_1 - T_2) = \frac{nR(T_1 - T_2)}{\gamma - 1},$$

which is exactly form 2. So forms 2 and 3 are also identical for any ideal gas.

All three forms are therefore one formula in three different costumes. Use whichever costume matches the given data: form 1 for pressures and volumes, form 2 for moles and temperatures, form 3 for an internal-energy cross-check. A well-set adiabatic problem will let you compute the answer with one form and verify it with another, which is the surest defence against arithmetic errors involving γ.

Why the adiabat is steeper than the isotherm

Both curves pass through A on the figure. From A, the isotherm follows PV = constant, which has slope dP/dV = −P/V along the curve. The adiabat follows PV^γ = constant, which has slope dP/dV = −γP/V. The adiabat's slope is steeper than the isotherm's by exactly the factor γ at every point. For γ = 1.4 (diatomic), the adiabat falls 1.4 times as fast per unit of V; for γ = 5/3 (monatomic), it falls 5/3 times as fast.

The physical reason for the steepness difference is the temperature. Isothermal compression keeps T fixed by allowing heat to leave. Adiabatic compression has no such heat outlet, so the temperature rises during the compression. A higher T means a higher P at any given V (because PV = nRT), so the adiabat reaches higher pressures for the same volume reduction than the isotherm does. Conversely, isothermal expansion absorbs heat to keep T fixed; adiabatic expansion has no heat to absorb, so T falls, and the adiabat reaches lower pressures for the same volume increase. Both effects make the adiabat steeper.

The adiabatic fingerprint

Five results identify every adiabatic process of an ideal gas. Each one will be reached for in the worked examples:

$\mathbf{Q = 0}$. The defining condition. No heat exchanged with the surroundings.

$\mathbf{PV^\gamma = constant.}$ The pressure-volume relation. The PV diagram traces an adiabat, steeper than an isotherm by the factor γ.

$\mathbf{TV^{\gamma-1} = constant.}$ Temperature-volume relation. The form to use when V is known and T is sought.

$\frac{T_2}{T_1} = \left(\frac{P_2}{P_1}\right)^{\frac{(\gamma-1)}{\gamma}}$. Temperature-pressure relation. The form to use when P is known and T is sought.

$\mathbf{\Delta U = -W = nC_v(T_1 - T_2).}$ The first law with Q = 0. Three equivalent forms of W are available, all reconcilable.

The walls insulate, the gas heats and cools by its own arithmetic, six examples now follow.

BINDER Example 27

0.15mol of a monatomic ideal gas, initially at 250kPa and 320K, expands adiabatically until its pressure drops to 100kPa. Find the final temperature of the gas and the work done by the gas during the expansion. Take R = 8.314Jmol^{-1}K^{-1}.

Solution

Monatomic means γ = 5/3.

Final temperature, from the temperature-pressure form of the adiabatic equation:

$$\frac{T_2}{T_1} = \left(\frac{P_2}{P_1}\right)^{\frac{(\gamma-1)}{\gamma}}$$

With $\frac{(\gamma-1)}{\gamma} = \frac{5/3-1}{5/3} = 0.4$:

$$T_2 = (320\text{K}) \times \left(\frac{100\text{kPa}}{250\text{kPa}}\right)^{0.4} = 222\text{K}$$

So the final temperature is 222K.
The gas has cooled by nearly one hundred degrees during its adiabatic expansion.
Work done by the gas, using the temperature form (form 2):

$$W = \frac{nR(T_1 - T_2)}{\gamma - 1} = \frac{(0.15\text{ mol})(8.314\text{Jmol}^{-1}\text{K}^{-1})(320\text{K} - 222\text{K})}{5/3 - 1} = 183\text{J}$$

Making Sense of the Answer: *A roughly hundred-kelvin cooling and 183J of work for a 0.15mol sample is a quietly impressive amount of physics from a small quantity of gas. The work done came out of the gas's own internal energy, so we can verify the result with form 3: $\Delta U = nC_v\Delta T = (0.15\text{mol})(3/2)(8.314\text{ J mol}^{-1}\text{ K}^{-1})(222\text{ K} - 320\text{ K}) \approx -183\text{J}$, and $W = -\Delta U = +183\text{J}$, agreeing exactly. Two completely different formulae, two routes through the algebra, one identical answer. That is the kind of internal consistency every adiabatic problem should carry.*

Think Like a Physicist: *The exponent $(\gamma-1)/\gamma$ is small for any real gas: 2/5 for monatomic, 2/7 for diatomic, 1/4 for polyatomic. This means the temperature ratio T_2/T_1 never changes by very much for moderate pressure ratios, even though the pressure ratio itself can be large. In this problem, the pressure ratio was 0.4 but the temperature ratio came out to 0.69, which is a much milder change. This is why adiabatic cooling, while real, is rarely catastrophic for gases that change pressure by factors of two or three. Order-of-magnitude pressure changes are needed to produce order-of-magnitude temperature changes adiabatically.*

BINDER Example 28

A motorcar tyre is at 4.0atm and 27°C. The tyre suddenly bursts, and the air inside expands to atmospheric pressure (1.0atm). Taking $\gamma = 1.4$ for the air, calculate the temperature of the escaping air immediately after the burst.

Solution

The expansion is **sudden**, fast enough that no heat has time to flow into the air through the surrounding tyre or atmosphere during the brief expansion event. So the process is adiabatic, and the temperature-pressure form of the adiabatic equation applies.

$$T_1 = (27 + 273)\text{K} = 300\text{K}$$

Apply the adiabatic temperature-pressure relation:

$$\frac{T_2}{T_1} = \left(\frac{P_2}{P_1}\right)^{\frac{(\gamma-1)}{\gamma}}$$

With $\frac{(\gamma-1)}{\gamma} = \frac{1.4-1}{1.4} = 0.286$:

$$T_2 = (300\text{K}) \times \left(\frac{1.0\text{atm}}{4.0\text{atm}}\right)^{0.286} = 202\text{K} = -71°\text{C}.$$

The air emerging from the burst tyre is therefore at about −71°C (almost a hundred degrees below the road surface temperature).

Making Sense of the Answer: *Negative seventy-one degrees Celsius sounds dramatic and is dramatic. But the cooling is brief: the cold air mixes within seconds with the surrounding atmospheric air, which is around 300K, and the temperature equilibrates rapidly. The −71°C is the temperature of the air for the very short interval immediately after the burst, before that mixing has had a chance to happen. It is what a thermometer pointed straight at the burst hole, in the first fraction of a second, would read. After a few seconds the cold air is gone and the residual air at the burst site is back to ambient. Adiabatic cooling is a transient effect, sharp but short.*

Think Like a Physicist: *This is the canonical exam-style adiabatic question: given the initial state and the final pressure, find the final temperature. The adiabatic temperature-pressure formula is the workhorse, and the only arithmetic challenge is the non-integer exponent. For $\gamma = 1.4$ (most exam problems), the exponent is $2/7 \approx 0.286$. For $\gamma = 5/3$ (monatomic), it is $2/5 = 0.4$. For $\gamma = 4/3$ (polyatomic), it is $1/4 = 0.25$. The rest of the calculation is direct substitution.*

BINDER Example 29

Air at atmospheric pressure (1.0×10^5Pa) is compressed adiabatically to half its initial volume. Taking $\gamma = 1.4$ for the air, find the final pressure.

Solution

From the adiabatic pressure-volume relation:

$$P_1 V_1^{\gamma} = P_2 V_2^{\gamma},$$

rearranging for P_2:

$$P_2 = P_1 \times \left(\frac{V_1}{V_2}\right)^{\gamma}$$

With $V_2 = V_1/2$, the volume ratio $V_1/V_2 = 2$:

$$P_2 = (1.0 \times 10^5\ \text{Pa}) \times (2)^{1.4} = 2.64 \times 10^5 \text{Pa}$$

So the final pressure is about 2.64×10^5 Pa.

Making Sense of the Answer: *Halving the volume isothermally would have doubled the pressure, by Boyle's law. Halving it adiabatically produces a* $2.64-$*fold increase, more than doubling. The extra factor of* 1.32 *comes from the temperature rise that accompanies adiabatic compression: the gas heats up, and that temperature rise contributes its own pressure increase on top of the geometric one from compression alone. The factor by which the adiabatic pressure exceeds the isothermal pressure for the same volume change is exactly* $\left(\frac{V_1}{V_2}\right)^{\gamma-1} = 2^{0.4} = 1.32$*, which matches:* $2.64 / 2.0 = 1.32$.

Think Like a Physicist: *The adiabatic compression ratio* $2^{\gamma} = 2.64$ *sits between two simpler benchmarks. Boyle's law for the same volume change would give* $2^1 = 2$ *(pure compression, no temperature effect). Real gases have γ between 1 and 5/3, so the adiabatic compression ratio for halved volume sits between 2.0 and* $2^{1.67} \approx 3.17$*. Most real gases come out at 2.6 ± 0.3, which is exactly where this answer sits.*

REAL Example 30

A scuba diver in Lake Tanganyika is at a depth of 30m, where the water pressure is approximately 4.0atm. She releases an air bubble from her regulator. The bubble, initially at the local water temperature of 22°C, rises rapidly to the lake's surface, where the pressure is just 1.0atm. Taking γ = 1.4 for the air inside the bubble, find the temperature of the air inside the bubble when it arrives at the surface, and explain in physical terms why the bubble is colder than the water it has risen through.

Solution

The word "rapidly" is the physical cue. The bubble's centimetre-scale interior with its thin gas-water boundary cannot conduct heat fast enough to keep up with a fast rise of several seconds, so during the rise **no** significant heat flows from the surrounding water into the air inside the bubble. With Q ≈ 0, the bubble's expansion is adiabatic.

$$T_1 = (22 + 273)\text{K} = 295\text{K}$$

With $(\gamma-1)/\gamma = 2/7 \approx 0.286$:

$$T_2 = (295\ \text{K}) \times \left(\frac{1.0}{4.0}\right)^{0.286} = 198\text{K} = -75°\text{C}.$$

So the air inside the bubble at the moment it breaks the surface is at about −75°C.

Physical explanation: As the bubble rises, the water pressure pushing inward on it decreases, so the bubble expands. The expanding air pushes back the water around it, doing work on the water; work that comes out of the air's internal energy because the rise is fast and there is no time for heat to flow into the bubble from the surrounding water through the bubble's thin gas-water interface. Falling internal energy means falling temperature. The faster the bubble rises, the closer this picture is to perfectly adiabatic. A slow rise would allow the warm water to keep the bubble at ambient temperature throughout (an isothermal expansion). The bubble's actual behaviour is somewhere between, but for a typical fast-rising bubble it is close enough to adiabatic that the calculation above gives a useful estimate.

Making Sense of the Answer: *−75°C is sharp cooling for a bubble rising through 22°C water. In practice, the bubble does not stay at −75°C for long: as it breaks the surface, the air inside mixes with the atmosphere and warms back to ambient within seconds. The −75°C is the air's transient state during the rise itself, and is what would be measured by a tiny thermometer riding inside the bubble. The bubble's wall, and the water immediately surrounding the bubble, would also cool slightly through contact with the cold air, but this transient effect is too small to be felt by the diver.*

Think Like a Physicist: *The adiabatic model applies whenever the timescale of the volume change is short compared with the timescale of heat conduction through the boundary. For a centimetre-scale bubble rising metres per second through water, the rise takes a few seconds; heat conduction across the bubble's tiny diameter takes considerably longer. The adiabatic approximation is therefore reasonable. For a slow geological process such as magma rising through kilometres of crust over millions of years, the expansion is effectively isothermal, even though the volumes are vastly larger and the pressure changes far more dramatic. Speed and length scales together decide which idealization applies; the underlying physics does not change with size, only with the ratio of process time to heat-flow time.*

HOT Example 31

An ideal gas (n = 2.0mol, γ = 1.4) is initially at 300K, 1.0×10^5Pa, and occupies a volume of 0.050m³. The gas is compressed adiabatically until its volume is reduced to 0.020m³. Find: (a) the final temperature, (b) the final pressure. (c) Verify the two answers are consistent with the ideal-gas law. (d) Calculate the work done on the gas during the compression. Take R = $8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) Final temperature from $TV^{\gamma-1}$ = constant:

$$T_2 = T_1 \times \left(\frac{V_1}{V_2}\right)^{\gamma-1}$$

With $\gamma - 1 = 0.4$:

$$T_2 = (300\text{K}) \times \left(\frac{0.050\text{m}^3}{0.020\text{m}^3}\right)^{0.4} = 433\text{K}$$

(b) Final pressure from PV^{γ} = constant:

$$P_2 = P_1 \times \left(\frac{V_1}{V_2}\right)^{\gamma} = (1.0 \times 10^5\ \text{Pa}) \times (2.5)^{1.4} = 3.6 \times 10^5\ \text{Pa}$$

(c) At state 2:

$$P_2 V_2 = (3.609 \times 10^5\ \text{Pa})(0.020\ \text{m}^3) = 7218\text{J},$$

$$nRT_2 = (2.0\ \text{mol})(8.314\ \text{J mol}^{-1}\text{K}^{-1})(433\ \text{K}) = 7200\text{J} \approx 7218\text{J}\ (P_2 V_2)$$

These agree to within rounding ($P_2 V_2 \approx nRT_2$), as the ideal-gas law requires. The adiabatic answers are internally consistent.

(d) Work done **by** the gas, from form 2:

$$W_{\text{gas}} = \frac{nR(T_1 - T_2)}{\gamma - 1} = \frac{(2.0\ \text{mol})(8.314\ \text{J mol}^{-1}\ \text{K}^{-1})(300\ \text{K} - 433\ \text{K})}{0.4} = -5530\text{J}$$

So the work done by the gas is −5530J, meaning the work done on the gas by the surroundings is +5530J:

$$W_{\text{on gas}} = +5530\ \text{J} = 5.5\text{kJ}$$

Verification using the internal-energy form. The change in internal energy is:

$$\Delta U = nC_v\Delta T = (2.0\ \text{mol}) \times \frac{R}{\gamma - 1} \times (T_2 - T_1)$$

$$= (2.0\ \text{mol}) \times \frac{8.314\ \text{J mol}^{-1}\ \text{K}^{-1}}{0.4} \times (433\text{K} - 300\text{K}) = +5530\ \text{J}$$

Since Q = 0, ΔU equals the work done on the gas, and both come out to 5.5kJ. Two completely different formulae deliver the same answer, as the first law of thermodynamics requires for any adiabatic process.

Making Sense of the Answer: *Five and a half kilojoules of work compresses 2 moles of gas from 50L to 20L, raising its temperature from room temperature to 160°C in the process. All 5.5kJ of that work ends up stored as internal energy in the gas, because Q = 0 means no energy leaks out as heat. This is exactly why a tightly compressed gas cylinder, freshly filled, feels warm to the touch, the compressor's work has gone into the gas. The two-route consistency check (PV-and-V-form work versus internal-energy form) confirms the answer in two independent ways, which is the safest possible check for a calculation involving exponents and γ.*

Think Like a Physicist: *The four parts of this problem map directly onto the four ways to think about an adiabatic process. Part (a) used the temperature-volume relation; part (b) used the pressure-volume relation; part (c) used the ideal-gas law as a consistency check between the two; part (d) used the work formula and confirmed it via the internal-energy form. Each part is a different face of the same underlying physics. Being able to switch fluidly between these four perspectives, and using one to check another, is the mark of a student who has mastered adiabatic processes rather than just memorised a formula. Every adiabatic question can be checked this way.*

HOT Example 32

One mole of oxygen gas ($\gamma = 1.4$) is initially at 2.0atm and 320K, occupying some volume V. The gas is first compressed isothermally until its volume halves. Then, starting from the end state of the first compression, the gas expands adiabatically and reversibly back to its original volume V. Find:

(a) The pressure of the gas at the end of the isothermal compression.

(b) The work done by the gas during the isothermal compression.

(c) The pressure of the gas at the end of the adiabatic expansion.

(d) The temperature of the gas at the end of the adiabatic expansion.

(e) Comment on whether the gas has returned to its initial state, and explain in physical terms.

Solution

Three states are involved in this problem, and naming them clearly at the start prevents confusion later. **State 1** *is the initial state.* **State 2** *is the end of the isothermal compression.* **State 3** *is the end of the adiabatic expansion.*

State 1: $P_1 = 2.0$atm, $T_1 = 320$K, volume = V (a quantity left as a symbol because no numerical V is given).

(a) Isothermal compression from V to V/2. Boyle's law applies because the temperature is held constant:

$$P_1V_1 = P_2V_2,$$

$$P_2 = P_1 \times \frac{V_1}{V_2} = (2.0\ \text{atm}) \times \frac{V}{V/2} = (2.0\text{atm}) \times 2 = 4.0\text{atm}$$

The pressure of the gas at the end of the isothermal compression is 4.0atm.

(b) Work done by the gas during the isothermal compression:

$$W_{12} = nRT \ln \frac{V_2}{V_1} = (1.0\ \text{mol})(8.314\ \text{J mol}^{-1}\ \text{K}^{-1})(320\ \text{K}) \times \ln \frac{V/2}{V} = -1844\text{J} \approx -1.84\text{kJ}$$

The work done by the gas is -1844J (or -1.84kJ)

The negative sign means the surroundings did 1.84kJ of work on the gas during the compression. The gas itself did −1.84kJ of work on the surroundings.

(c) Adiabatic expansion from $V_2 = V/2$ back to $V_3 = V$. The starting state of this expansion is **State 2**:

$P_2 = 4.0$atm, $V_2 = V/2$, $T_2 = 320$K. Using the adiabatic pressure-volume relation:

$$P_2V_2^{\gamma} = P_3V_3^{\gamma},$$

$$P_3 = P_2 \times \left(\frac{V_2}{V_3}\right)^{\gamma} = (4.0\ \text{atm}) \times (V/2 \div V)^{1.4} = (4.0\ \text{atm}) \times (0.5)^{1.4} = 1.52\text{atm}$$

The pressure of the gas at the end of the adiabatic expansion is 1.52atm.

(d) Temperature at State 3, using the adiabatic temperature-volume relation:

$$T_2V_2^{\gamma-1} = T_3V_3^{\gamma-1},$$

$$T_3 = T_2 \times \left(\frac{V_2}{V_3}\right)^{\gamma-1} = (320\ \text{K}) \times (0.5)^{0.4} = 243\text{K}$$

The temperature of the gas at the end of the adiabatic expansion is 243K.

(e) Comment on the final state: The gas has returned to its original volume V, but it has **not** returned to its original state. Three numbers tell the story:

Volume: $V_3 = V = V_1$. **Equal** to the initial value.

Pressure: $P_3 = 1.52$atm, less than $P_1 = 2.0$atm. **Lower** than initial.

Temperature: $T_3 = 243$K, less than $T_1 = 320$K. **Lower** than initial.

Why has the gas not returned to its original state, despite returning to its original volume? Because the two processes were genuinely different and did not cancel out. The isothermal compression took heat out of the gas ($Q = W = -1.84$kJ) into the surroundings while leaving the temperature unchanged. The adiabatic expansion that followed exchanged no heat at all with the surroundings, but lost work-energy from inside the gas, so the temperature fell. The gas ended at a state with lower T, lower P, and the original V. To return the gas exactly to State 1 from State 3, an additional process step would be required: heating the gas at constant volume from 243K back to 320K, which would also restore the pressure (because PV = nRT and V is fixed). Different paths between the same volumes can leave the gas at very different states.

Making Sense of the Answer: *The volume returned to its starting value, but the gas's other state variables did not. Of the 1.84kJ of work the surroundings put into the gas during the isothermal compression, the gas immediately dumped 1.84kJ as heat to the heat reservoir; nothing was stored. Then in the adiabatic expansion, the gas spent some of its own internal energy doing work against the surroundings, and the temperature fell as a result. The two processes have completely different energy bookkeeping: the isothermal one trades work in for heat out and conserves U; the adiabatic one trades U for work and conserves Q. Doing them in sequence does not undo the first with the second. State functions (V) returned to their starting values; path functions (Q and W along each step) did not, and the cumulative path leaves the gas elsewhere on the PV diagram.*

Think Like a Physicist: *The PV diagram is the most efficient way to see why this problem ends where it does. The isothermal step traces a hyperbola from State 1 to State 2 (going leftward, up and to the left, since volume halves and pressure doubles). The adiabatic step traces a steeper curve from State 2 back to volume V, but because the adiabat is steeper than the isotherm, it does not return to the same starting pressure when it reaches the original volume. It lands lower, at State 3 directly below State 1 on the diagram. The gap between State 1 and State 3, both at volume V, is exactly the difference between an isotherm and an adiabat through State 2. Examination questions often test this kind of cyclic-but-not-quite reasoning, because it forces the student to recognise that path-dependence applies to compound processes too: not just W and Q for one process, but also the entire net thermodynamic state at the end of a sequence.*

Four thermodynamic processes are now in hand: V fixed, P fixed, T fixed, and Q fixed. Each one pinned down a different state variable and let the others adjust; each one carries its own work formula and its own characteristic curve on the PV diagram. The next section assembles all four into a single mastery picture, with a master summary table that the student should learn to recognise and apply, and a synthesis figure showing all four processes radiating from a common starting point. The combined-process example just delivered is one species of multi-step problem; the next section systematises the technique for handling them.

MASTERING THE PV DIAGRAM

The four cornerstone thermodynamic processes are now in hand. Each one was met as a separate problem in its own section, with its own picture, its own derivation, and its own worked examples. This closing section pulls the four together. The aim is twofold. First, to assemble what has been learned into a single mastery picture and a single summary table that can be revisited at any moment to verify a half-remembered formula. Second, to graduate from individual processes to the multi-process problems that examiners actually set, in which the gas walks through two or three processes in sequence and the student must account for the total work and total heat across all of them.

Imagine a gas at some particular state, with a particular pressure, volume, and temperature. From that single starting point, the gas can be carried in any of four canonical directions, each fixing one of the variables. The picture below shows all four processes radiating from a common starting state A on the PV diagram. This is the master diagram of the chapter, worth keeping near to hand whenever a thermodynamics problem is being read.

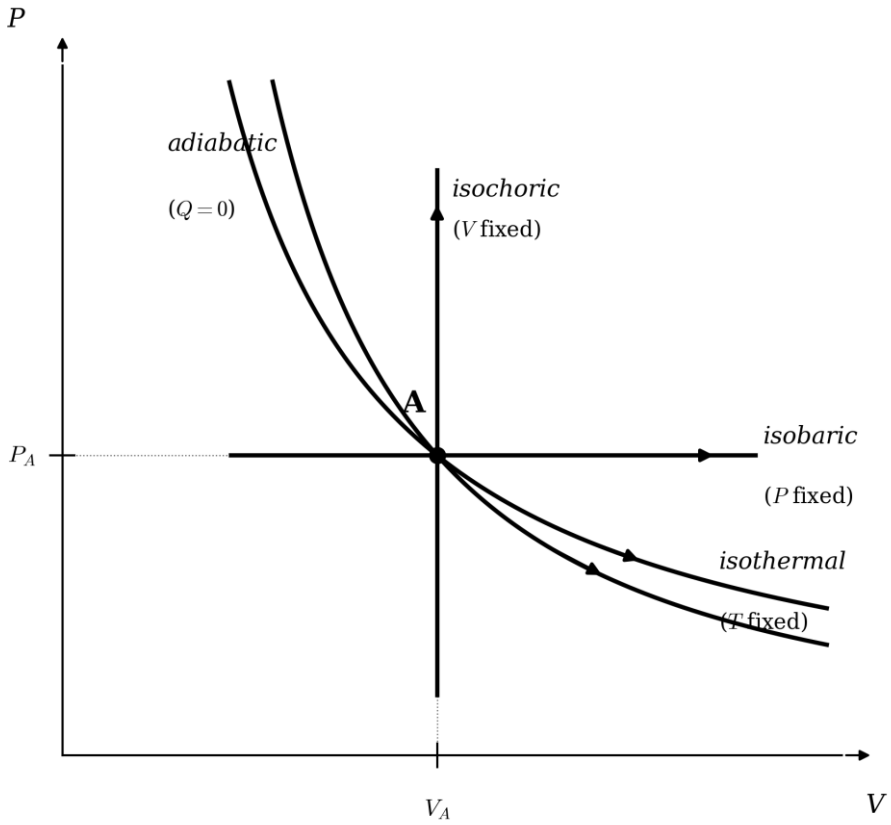


Figure: ***The four cornerstone thermodynamic processes drawn from a common starting point A on a PV diagram.*** *The vertical line is the isochoric process (V fixed, only P and T change). The horizontal line is the isobaric process (P fixed, only V and T change). The gentler hyperbola is the isothermal process (T fixed, PV = constant). The steeper hyperbola is the adiabatic process ($Q = 0, PV^{\gamma} = \text{constant}$). At point A all four processes pass through the same state; from A onward they diverge. Knowing which way the gas's state moves on this diagram tells the student immediately which formula to reach for.*

The master process summary table

The four formulae for Q, W, and ΔU across the four processes can be assembled into a single comparison table. Every entry is derivable from what was established in the four preceding sections; the table simply puts them side by side so the patterns become visible at a glance.

Process	Constant	Process equation	Q	W	ΔU
Isochoric	V	$\frac{P}{T} = \text{const}$	$nC_v\Delta T$	0	$nC_v\Delta T$
Isobaric	P	$\frac{V}{T} = \text{const}$	$nC_p\Delta T$	$P\Delta V = nR\Delta T$	$nC_v\Delta T$
Isothermal	T	$PV = \text{const}$	$= W$	$nRT\ln\frac{V_2}{V_1}$	0
Adiabatic	Q = 0	$PV^{\gamma} = \text{const}$	0	$\frac{P_1V_1 - P_2V_2}{\gamma - 1}$	$= -W$

Three patterns deserve highlighting.

- Two of the four processes have one zero entry each, but in different columns. Isochoric has W = 0, because no volume change means no work. Adiabatic has Q = 0, by definition. The remaining quantity in each case is computed from the first law: for the isochoric process, all the heat goes into internal energy, so Q = ΔU; for the adiabatic process, all the work comes out of internal energy, so W = −ΔU.
- Although the ΔU column shows different expressions for each process, the formula $\Delta U = nC_v\Delta T$ applies to every process of an ideal gas, because the internal energy depends only on temperature. The table simply chooses the most illuminating form for each row: $nC_v\Delta T$ for isochoric and isobaric, where the formula is the natural way to compute ΔU; zero for isothermal, which is just $nC_v\Delta T$ with ΔT = 0; and −W for adiabatic, which is the first law with Q = 0. Whether the volume or pressure changes is irrelevant for ΔU; only ΔT matters, and ΔT might be zero or might be derivable from W.
- The entries in the Q column simplify dramatically for two of the four. Isothermal Q equals W, and adiabatic Q is zero by definition. Each of these simplifications has been derived in its own section and should now be remembered as a feature of the table.

Closure of cycles

A particularly important application of the table is to thermodynamic cycles. A cycle is a sequence of processes that returns the gas to exactly the state at which it started. State variables P, V, T, and U are by definition unchanged at the end of a cycle, since the gas is back where it began. So:

$$\Delta U_{\text{cycle}} = 0,$$

regardless of how complicated the path was. The first law of dynamics for a complete cycle then gives:

$$Q_{\text{cycle}} = W_{\text{cycle}}$$

Net heat absorbed by the gas over the cycle equals net work done by the gas over the cycle. This is how every heat engine in the world operates: heat absorbed at one stage of the cycle, work delivered at another, and the gas returned to the start to repeat the procedure. The petrol engine in a car, the diesel engine in a daladala, and the gas turbine at Ubungo power station in Dar es Salaam all exploit exactly this closure. The numerical net work per cycle equals the area enclosed by the cycle on the PV diagram, traversed in its appropriate direction. A clockwise cycle encloses a positive area, meaning W is positive: the gas does net work on its surroundings, and the device is producing power. A counter-clockwise cycle encloses a negative area, meaning W is negative: the surroundings do net work on the gas, and the device is consuming power. Engines are clockwise; refrigerators and heat pumps are counter-clockwise. The direction of traversal alone tells us which.

Combined processes and the state-transfer rule

Many examination problems describe a sequence of two or more thermodynamic processes carried out in succession. A gas might be compressed isothermally and then heated isochorically, or expanded adiabatically and then cooled isobarically. The general principle for handling such combined-process problems is short and important enough to deserve its own statement.

The state-transfer rule: *In any sequence of thermodynamic processes, the final state of one process is the initial state of the next. The values of P, V, T, and n at the end of process 1 carry over unchanged as the starting values for process 2.* The numbers do not reset; they propagate forward.

This means a two-process problem is solved by treating the two processes in turn. Use the formulae appropriate to process 1 to find its final state. Take that final state as the initial state of process 2. Use the formulae appropriate to process 2 to find its final state. Sum any quantities that accumulate (work, heat) across the two processes; carry forward state variables (P, V, T, n) in their final values.

A common student error is to apply the initial conditions of process 1 to the start of process 2, effectively rewinding the gas to before process 1 took place. This must not be done. The state-transfer rule is the principle that prevents the error, and it must be invoked explicitly in every combined-process problem.

Real-world applications of the four processes

Real machines and natural processes rarely involve a single thermodynamic process in isolation. Most consist of sequences of processes in cycles, or approximations to single processes that can be analysed using the appropriate idealization. A few examples worth knowing.

1. Petrol engine (four-stroke)

The four strokes of a petrol engine map one-to-one onto the four cornerstone processes, in approximate form.

- **Intake stroke:** the piston descends and air-fuel mixture is drawn in at roughly atmospheric pressure (isobaric).
- **Compression stroke:** the piston rises rapidly, compressing the mixture before significant heat can leak through the cylinder walls (nearly adiabatic).
- **Combustion:** the spark plug fires and the fuel burns explosively while the piston is still near top dead centre, so the combustion happens at nearly constant volume (nearly isochoric).

- **Power stroke:** the hot high-pressure gas pushes the piston back down rapidly, expanding nearly adiabatically.
- **Exhaust stroke:** the burnt gas is pushed out at near-atmospheric pressure (isobaric).

All four cornerstone processes appear in a single cycle of the engine, in a sequence that has been refined and tuned by mechanical engineers for more than a century but whose underlying physics is exactly the four cases this chapter has built up.

2. Diesel engine

Differs from the petrol engine in one critical respect: there is no spark plug. The compression ratio is so high that the adiabatic compression of the air alone raises its temperature above the ignition point of diesel fuel. Fuel is then injected directly into the hot compressed air and ignites on contact, with no spark needed. The temperature rise during compression is computed using $TV^{\gamma-1}$ = constant, which is a standard adiabatic calculation, and is the reason a diesel engine has its higher compression ratio than a petrol engine.

3. Refrigerator and air conditioner

These devices run a thermodynamic cycle in the opposite direction to an engine. A working fluid (typically a refrigerant gas) is compressed adiabatically by a motor-driven compressor, raising its temperature. The hot fluid then releases heat to the kitchen or outside air at near-constant pressure (isobaric). The fluid is then expanded back through a valve, cooling adiabatically. The cold fluid then absorbs heat from inside the refrigerator at near-constant pressure (isobaric). The cycle is counter-clockwise on the PV diagram, the work done on the gas is positive, and the device is a refrigerator (consuming electrical energy to move heat from the cold compartment to the warm room) rather than an engine (producing energy from a temperature difference).

4. Speed of sound in air

Sound waves are alternating compressions and rarefactions of air, with frequencies in the hundreds or thousands of hertz. At such high frequencies, each compression or rarefaction lasts only milliseconds, far too short for heat to flow significantly through the air during one cycle. Sound propagation in air is therefore an adiabatic, not isothermal, phenomenon, with the speed of sound given by the formula, $\sqrt{(\gamma RT/M)}$, where γ enters because the propagation is adiabatic. The fact that γ rather than 1 appears in this formula was first appreciated by Laplace, correcting an earlier calculation by Newton that had assumed isothermal propagation and produced a speed about fifteen percent too low.

5. Gas turbine power station

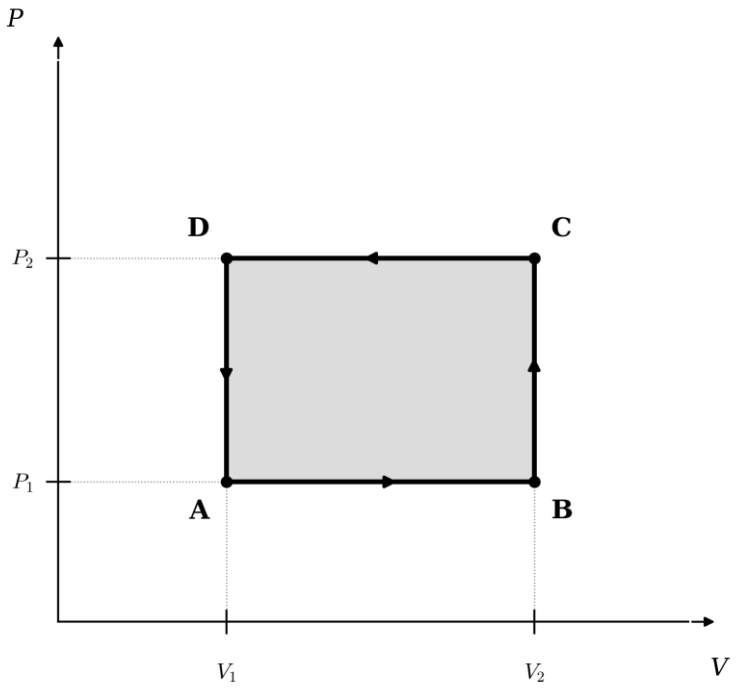
The gas turbines that supply Dar es Salaam with much of its electrical power, and the diesel and gas-fired backup units at Kidatu and elsewhere, exploit a Brayton-style cycle. Air is compressed in a multi-stage compressor (nearly adiabatic), heated in a combustion chamber by burning natural gas at constant pressure (isobaric), expanded through the turbine blades to drive the generator (nearly adiabatic), and finally exhausted to the atmosphere (isobaric). The cycle, drawn on a PV diagram, is a closed loop traversed clockwise, and its enclosed area is the net work delivered to the generator per cycle. The same physics that runs in a daladala engine runs, scaled up by orders of magnitude, in the turbine halls of every power station in the country.

The four processes have now been brought together, three examples now follow.

BINDER Example 33

A fixed amount of an ideal gas is taken around a rectangular cycle on a PV diagram, as shown in the figure. The corners of the rectangle are at A(V_1, P_1), B(V_2, P_1), C(V_2, P_2), and D(V_1, P_2), with $V_1 = 1.0\times10^{-3}\text{m}^3$, $V_2 = 3.0\times10^{-3}\text{m}^3$, $P_1 = 1.0\times10^5\text{Pa}$, and $P_2 = 2.0\times10^5\text{Pa}$. The gas is taken around the cycle in the order A→B→C→D→A.

(a) Identify each leg of the cycle by the type of thermodynamic process it represents.
(b) Compute the net work done by the gas per cycle.
(c) State whether the net heat flow over the cycle is into or out of the gas.



Solution

(a) Identification of each leg:

A→B. The gas moves from (V_1, P_1) to (V_2, P_1), so the volume increases from V_1 to V_2 while the pressure stays fixed at P_1. This is an **isobaric expansion** at low pressure.

B→C. The gas moves from (V_2, P_1) to (V_2, P_2), so the pressure increases from P_1 to P_2 while the volume stays fixed at V_2. This is an **isochoric pressure rise** at large volume.

C→D. The gas moves from (V_2, P_2) to (V_1, P_2), so the volume decreases from V_2 to V_1 while the pressure stays fixed at P_2. This is an **isobaric compression** at high pressure.

D→A. The gas moves from (V_1, P_2) to (V_1, P_1), so the pressure decreases from P_2 to P_1 while the volume stays fixed at V_1. This is an **isochoric pressure drop** at small volume.

(b) The two isochoric legs (B→C and D→A) contribute zero work, because the volume does not change during them. Only the two isobaric legs do work.

On A→B, the gas expands at the low pressure P_1:

$$W_{A\to B} = P_1(V_2 - V_1) = (1.0\times10^5\text{Pa})(3.0\times10^{-3}\text{m}^3 - 1.0\times10^{-3}\text{m}^3) = +200\text{J}$$

On C→D, the gas is compressed at the high pressure P_2:

$$W_{C\to D} = P_2(V_1 - V_2) = (2.0\times10^5\text{Pa})(1.0\times10^{-3}\text{m}^3 - 3.0\times10^{-3}\text{m}^3) = -400\text{J}$$

The net work done by the gas per cycle is the sum of the four legs:

$$W_{cycle} = W_{A\to B} + W_{B\to C} + W_{C\to D} + W_{D\to A}$$
$$= +200\text{J} + 0\text{J} + (-400\text{J}) + 0\text{J} = -200\text{J}$$

So the gas does −200J of net work per cycle.

The negative sign tells us that the surroundings have done 200J of net work on the gas over the complete cycle, not the other way round.

Cross-checking with the enclosed-area method: the rectangle has width $V_2 - V_1 = 2.0\times10^{-3}\text{m}^3$ and height $P_2 - P_1 = 1.0\times10^5\text{Pa}$, so its area is $(2.0\times10^{-3}\text{m}^3)(1.0\times10^5\text{Pa}) = 200\text{J}$. The magnitude matches. The sign is supplied by the direction of traversal. Tracing the arrows on the figure, the path A→B→C→D→A goes rightward at the bottom, up on the right, leftward at the top, down on the left. This is counter-clockwise on a standard PV diagram. Counter-clockwise traversal makes the net work negative, which agrees with the leg-by-leg calculation: $W_{cycle} = -200\text{J}$.

(c) By cycle closure, $Q_{cycle} = W_{cycle} = -200\text{J}$. Negative net heat means heat flows out of the gas over the complete cycle.

The gas absorbs heat during some legs and rejects more heat during others; the net is outward. Since work flows in (the surroundings drive the gas around the cycle) and heat flows out (the gas dumps heat to its surroundings), this cycle is a refrigeration-style device. It uses external work to move heat from the gas to the environment.

Making Sense of the Answer: *Two-hundred joules of net heat exit the gas per cycle, and the same magnitude of work is supplied to the gas per cycle. This is the energy bookkeeping of a refrigeration-style device. The four legs do not all contribute work, and not all in the same direction; the net work is what survives after the four are summed. The leg-by-leg sum and the enclosed-area calculation agree on both the magnitude (200J) and the sign (negative, by the counter-clockwise traversal). The lesson: the sign of net work in a cycle problem is determined not by the shape of the cycle but by the direction in which the cycle is traversed. Direction matters as much as area.*

Think Like a Physicist: *Two methods are available for net work in a cycle. Method one: leg-by-leg calculation, summing the work done during each individual process and using the appropriate formula for that process. This method is foolproof but slow. Method two: enclosed-area calculation, which gives the magnitude in one step (just the area on the PV diagram), with the sign supplied by the direction of traversal (clockwise positive, counter-clockwise negative). The two methods must agree, and a careful student uses one to compute and the other to check. When they disagree, as in the discussion above, the disagreement points to an error in identifying the direction, which a careful re-examination of the arrows resolves.*

REAL Example 34

Kipute is studying a textbook diagram showing the working fluid of an air conditioner used in a Dar es Salaam supermarket cold-storage room, traced as a closed cycle on a PV diagram. The cycle is drawn in the counter-clockwise direction. Most engine cycles encountered in the chapter are drawn clockwise.

(a) Explain what the direction of traversal of a cycle on a PV diagram tells us about the net work done by the gas over the complete cycle.
(b) Explain why a heat engine cycle must be clockwise on a PV diagram and a refrigerator cycle must be counter-clockwise.
(c) State, with reason, the direction of net heat flow over the air conditioner's cycle.

Solution

(a) The net work done by the gas over a complete cycle equals the area enclosed by the cycle on the PV diagram, signed according to the direction of traversal. A clockwise cycle encloses a positive area; the net work done by the gas is positive, meaning the gas delivers work to its surroundings on balance. A counter-clockwise cycle encloses a negative area; the net work done by the gas is negative, meaning the surroundings do work on the gas on balance. The direction of traversal therefore determines the sign of the net work; the area determines its magnitude.

(b) A heat engine exists to produce work from a temperature difference. The gas inside the engine must, on balance, be doing work on the surroundings, since that work is the engine's useful output (turning a turbine blade or driving a piston). The net work done by the gas must therefore be positive, and from part (a) this requires the cycle to be traversed clockwise. A refrigerator or air conditioner does the opposite. It pumps heat from a cold compartment to a warmer surrounding, against the natural flow direction, by consuming work supplied from outside (typically electrical work driving a compressor). The work flows into the gas, so the net work done by the gas is negative, and the cycle must be traversed counter-clockwise. The direction of the loop on the PV diagram therefore identifies the device unambiguously: clockwise for an engine, counter-clockwise for a refrigerator or heat pump.

(c) For the air conditioner, the cycle is counter-clockwise, so the net work done by the gas is negative. By cycle closure (ΔU over a complete cycle is zero, so $Q_{cycle} = W_{cycle}$), the net heat flow over the cycle is also negative. Negative net heat means heat flows out of the gas on balance over the cycle, and into the surroundings. This is exactly what an air conditioner must do: it absorbs heat from inside the cold-storage room (where the gas is at low temperature during one part of the cycle), and rejects a larger amount of heat to the warmer outside air (where the gas is at high temperature during another part of the cycle). The difference between the heat rejected and the heat absorbed equals the work supplied by the electrical compressor, again by the first law applied over the cycle.

Making Sense of the Answer: *The supermarket air conditioner is correctly drawn counter-clockwise because that is what a refrigerator must look like on a PV diagram. The cycle's direction is its identity card: engines run clockwise, refrigerators run counter-clockwise. Both involve the same four cornerstone processes (in various combinations), the same first law, and the same kind of PV diagram, but with the loop traversed in opposite directions. The energy bookkeeping over the cycle differs only in sign, and the sign is read from the direction of the arrows.*

Think Like a Physicist: *Direction of traversal tells you the sign of the net work (clockwise positive, counter-clockwise negative). The sign of the net work tells you the device type (positive for engine, negative for refrigerator or heat pump). The sign of the net work also tells you the sign of the net heat flow, by cycle closure. Three pieces of physics, all read off a single arrow on a single diagram. When an examination question shows a PV cycle and asks for the device type or the direction of net heat flow, the answer is found by glancing at the direction of the arrows, not by computing anything. Reading direction first, numbers second, is the habit this section is most concerned to install.*

HOT Example 35

2500mol of air ($\gamma = 1.4$) is cooled at constant pressure of 1.0atm from 35°C to 26°C. Find (a) the molar heat capacity at constant volume of the gas (b) the change in internal energy of the gas, (c) the work done by the gas, and (d) the heat supplied to the gas. Take $R = 8.314 Jmol^{-1}K^{-1}$.

Solution

Note that the sample is cooled, not heated, so the temperature change is negative:

$$\Delta T = 26°C - 35°C = -9K$$

(a) The molar heat capacity at constant volume:

$$C_v = \frac{R}{\gamma - 1} = \frac{8.314 Jmol^{-1}K^{-1}}{1.4 - 1} = \frac{8.314 Jmol^{-1}K^{-1}}{0.4} = 20.79 Jmol^{-1}K^{-1}.$$

(b) The change in internal energy depends only on temperature for an ideal gas:

$$\Delta U = nC_v\Delta T = (2500mol)(20.79 Jmol^{-1}K^{-1})(-9K) = -467775J \approx -468kJ$$

Negative ΔU means the internal energy of the gas has fallen, which makes sense: the gas was cooled, so its molecular kinetic energy fell.

(c) The work done by the gas at constant pressure, using $W = nR\Delta T$:

$$W = nR\Delta T = (2500mol)(8.314 Jmol^{-1}K^{-1})(-9K) = -187065J \approx -187kJ$$

Negative W means the gas did not do positive work on its surroundings; instead, the surroundings did positive work on the gas as it contracted. This too makes sense: a cooling gas at constant pressure shrinks, and the atmosphere outside pushes it inward.

(d) The heat supplied to the gas, from the first law:

$$Q = \Delta U + W = (-468kJ) + (-187kJ) = -655kJ$$

Negative Q means heat did not flow into the gas; it flowed out of the gas. The cooling process delivered 655kJ of heat from the gas to the surroundings.

Verification using the alternative isobaric heat formula:

$$Q = nC_p\Delta T, \quad \text{where} \quad C_p = C_v + R = 20.79 Jmol^{-1}K^{-1} + 8.314 Jmol^{-1}K^{-1} = 29.10 Jmol^{-1}K^{-1}.$$

$$Q = (2500mol)(29.10 Jmol^{-1}K^{-1})(-9K) = -654750J \approx -655kJ$$

Both routes agree, as the first law requires.

Making Sense of the Answer: *Three negative numbers tell a single coherent story. The gas cooled (ΔT negative), so its internal energy fell (ΔU negative). The gas contracted as it cooled (because pressure was held constant, so V is proportional to T by Charles's law), so the work done by the gas was negative (W negative). And heat was extracted from the gas (Q negative). All three signs follow from the single physical fact: the gas was cooled. The energy bookkeeping closes: 468kJ of heat removed comes from the falling internal energy, and an additional 187kJ comes from the work done by the surroundings as they squeezed the contracting gas inward. Total heat extracted: 655kJ. The two-route check (using C_p directly versus going through C_v and the first law) confirms the arithmetic.*

Think Like a Physicist: *$nC_p\Delta T$ is the fastest route to Q for an isobaric process; $nC_v\Delta T$ plus $P\Delta V$ is the longer route via the first law. Both work, both should agree, and a careful student uses both: the first to compute, the second to check. The relation between the two is Mayer's equation, which is the reason C_p exceeds C_v by exactly R, and is the reason any isobaric problem can be cross-checked with this kind of two-route method. The 2500-mole sample size in this problem is realistic for a small industrial process; the magnitude of the heat (655kJ) sets the scale of how much cooling capacity the surrounding refrigeration system must deliver per such cooling event. Industrial thermodynamics works in exactly these scales.*

This closes the systematic study of the four cornerstone thermodynamic processes and their applications in cycles. The first law of thermodynamics has been the constant companion of every problem in this chapter: it has told us how heat, work, and internal energy must balance in every process, and it has imposed the cycle-closure condition that makes engines and refrigerators work the way they do. The first law accounts for energy. But energy accounting alone does not tell the whole story of thermodynamics. It does not, for instance, tell us why heat flows from hot bodies to cold bodies and not the other way; energy would be conserved either way. To understand the directionality of natural processes, a second law is needed, and the next section introduces it. The second law is what gives the difference between possible and impossible processes its precise mathematical meaning, and it is the foundation on which the concept of entropy is built.

THE SECOND LAW OF THERMODYNAMICS: A BRIEF INTRODUCTION

The first law of thermodynamics has accompanied every problem in this chapter so far, telling us how heat, work, and internal energy must balance in any process. The first law is a law of **conservation**: energy is neither created nor destroyed; it merely changes form. But energy conservation alone does not capture everything thermodynamics has to say. Two facts about the

natural world are not predicted, and not even forbidden, by the first law. **The first** is that *some processes happen spontaneously while their reverses, despite being equally allowed by energy conservation, never do.* **The second** is that *heat-engine efficiency has a ceiling that depends only on the temperatures of the hot and cold reservoirs, and not on the engineering skill of the designer*. To capture these facts, a second law of thermodynamics is needed. This closing section introduces that law qualitatively and applies it to the kind of efficiency calculations that examiners frequently set.

Place a hot cup of tea and a cold cup of water side by side on the classroom desk. Wait. After some minutes the tea has cooled, the water has warmed, and eventually both are at the same room temperature. Wait longer; nothing further happens. Now ask whether the reverse process could ever occur: a tepid cup at room temperature spontaneously dividing itself into a hotter portion and a colder portion, with no external agent acting on it. Energy conservation does not forbid such a thing. The total energy of the two cups before and after would be identical; only the distribution of that energy across the two cups would change. The first law would happily permit it. And yet no one in the history of physics has ever observed it. Hot bodies cool, cold bodies warm, and the equilibration runs in only one direction. The reverse, though energetically permitted, is forbidden by something else.

That something else is the **second law of thermodynamics**. It captures the directionality of natural processes, the fact that some directions in the space of physical possibilities are open and the reverse directions are closed. It is one of the most far-reaching statements in physics, and although the full mathematical machinery of the second law belongs to a more advanced course, the qualitative content of the law and its main quantitative consequence (the efficiency ceiling for heat engines) can be stated and used at this level.

Two classical statements of the second law

Two equivalent statements of the second law have come down through the history of thermodynamics. The first is due to Lord Kelvin and Max Planck, the second to Rudolf Clausius. Each captures the same underlying truth, but each starts from a different observation.

Kelvin-Planck statement: *No heat engine operating in a cycle can convert all of the absorbed heat into useful work. Some of the absorbed heat must always be rejected to a cooler body.*

This statement comes from observing real heat engines. The petrol engine in a car, the gas turbine in a power station, and any other heat-using device that runs in a cycle, all reject some of the heat they absorb to a cooler reservoir, in the form of warm exhaust, hot cooling water, or radiated heat. None of them converts all of the input heat into work. Engineers have tried, for two centuries, to design an engine that does. None has succeeded, and the second law tells us none ever will.

Clausius statement: *Heat cannot flow spontaneously from a cooler body to a hotter body. To move heat from cold to hot, work must be supplied from outside.*

This statement comes from observing the natural direction of heat flow. Heat always flows from hot to cold without external help; the reverse never occurs without external help. A refrigerator does indeed move heat from a cold compartment to a warmer kitchen, but only because an electric motor consumes work to drive the compressor. Unplug the refrigerator and the heat flow reverses, leaking from the warm kitchen back into the cold compartment, the compartment warming up to room temperature exactly as the second law predicts.

Equivalence: The Kelvin-Planck and Clausius statements look superficially different. One is about engines; the other is about heat flow. But they can be shown to be logically equivalent: any device that violates one of them could be combined with a normal device to violate the other. A full proof of this equivalence is beyond the scope of this book; for our purposes here, it is enough to know that the two statements are different verbal expressions of the same underlying physical principle. Both are correct, both are useful, and which to invoke in a given problem depends on whether the problem is about engines (use Kelvin-Planck) or about heat flow (use Clausius).

Carnot's theorem and the efficiency ceiling

The most useful quantitative consequence of the second law for examination purposes is **Carnot's theorem**, named for the French engineer Sadi Carnot, who established it about 200 years ago in a remarkable monograph written when he was just 28 years old. The theorem states a strict ceiling on the efficiency of any heat engine, expressed in terms of only **two quantities**: *the temperature of the hot reservoir from which the engine absorbs heat, **and** the temperature of the cold reservoir to which it rejects heat.*

We define the **efficiency** η of a heat engine in the natural way, as the ratio of useful work output per cycle to heat input per cycle:

$$\eta = \frac{\text{useful work output per cycle}}{\text{heat input per cycle}}$$

This efficiency lies between 0 (a useless engine that produces no work) and 1 (a perfect engine that converts all heat to work). The first law alone allows any value in this range. The second law restricts the range further. **Carnot's theorem** says:

No heat engine operating between a hot reservoir at absolute temperature T_h and a cold reservoir at absolute temperature T_c can have an efficiency greater than the Carnot efficiency:

$$\eta_{max} = 1 - \frac{T_c}{T_h},$$

where the temperatures must be expressed in kelvin. This bound is independent of the working substance, the design of the engine, and any other detail. It depends only on the two reservoir temperatures.

The ceiling has immediate practical consequences. Consider a typical petrol engine, in which the combustion temperature might reach about 800K and the surroundings sit at about 320K. The Carnot efficiency for these temperatures is:

$$\eta_{max} = 1 - \frac{320K}{800K} = 1 - 0.40 = 0.60 = 60\%$$

So the absolute ceiling on the efficiency of any engine running between these temperatures is 60%. No matter how cleverly the engine is designed, the second law forbids it from doing better. Real petrol engines achieve about 25% to 30% in practice, well below the ceiling, because of friction, incomplete combustion, heat loss through the cylinder walls, and other practical inefficiencies. But the second law tells us that even a perfectly designed engine, with no friction or losses, could not exceed 60% in this temperature range. The ceiling is not an engineering target; it is a law of nature.

This is also the long-promised explanation for why a working engine gets hot. If 100% efficiency were possible, an engine could absorb its fuel's chemical energy and deliver every joule as useful work, leaving no heat behind. Real engines must reject 70% to 75% of the absorbed heat to the surroundings, and that rejected heat is what warms the engine block, the radiator water, and the air around the exhaust. The mechanic at every petrol station in the country who shrugs at a hot engine and says "a hot engine is a working engine" is, knowingly or not, invoking the second law. The first law alone could not have told him so.

Two laws are now in hand, three examples now follow.

BINDER Example 36

Identify which form of the second law of thermodynamics is most directly invoked by each of the following statements:

(a) A refrigerator cannot work without a power supply.

(b) No heat engine can be 100% efficient.

Solution

(a) "A refrigerator cannot work without a power supply" is most directly an instance of the **Clausius form**. A refrigerator's job is to move heat from a colder compartment to a warmer kitchen, against the natural direction of heat flow. The Clausius form says exactly that this transfer requires work input from outside, which is the power supply mentioned in the statement.

(b) "No heat engine can be 100% efficient" is most directly an instance of the **Kelvin-Planck form**, which says exactly that some heat must always be rejected to a cooler reservoir, so the work output cannot equal the full heat input.

Making Sense of the Answer: *The two forms of the second law specialize to two different real-world problem types. Refrigerators, heat pumps, and any device that moves heat against the natural direction land under the Clausius form. Engines, turbines, and any device that converts heat to work land under the Kelvin-Planck form. The two forms are equivalent in the sense that any violation of one would imply a violation of the other, but in practice, choosing the form that matches the problem at hand makes the reasoning quicker.*

Think Like a Physicist: *Both classical statements of the second law are statements of impossibility. The first law is a positive statement (energy is conserved). The second law, by contrast, is a series of negative statements (this cannot happen, that cannot happen, no engine can be 100% efficient, no heat can flow spontaneously from cold to hot). Statements of impossibility are sometimes more powerful than statements of possibility, because every observed instance of the forbidden process not occurring is evidence for the law. After two centuries of careful observation, the second law has accumulated more evidence than almost any other law in physics, and despite the simplicity of its forms, no exception has ever been found.*

REAL Example 37

A Dar es Salaam taxi driver tells a passenger that a friend of his has designed an engine that recovers 95% of the fuel's energy as useful work. The engine operates between a combustion temperature of 700K and an ambient temperature of 305K.

(a) Use Carnot's theorem to determine whether such an engine is possible.
(b) Find the correct theoretical maximum efficiency for an engine operating between these two temperatures, and state in physical terms what the result implies about the friend's claim.

Solution

(a) Carnot's theorem sets the upper bound on the efficiency of any heat engine operating between two reservoir temperatures. Computing that bound for $T_h = 700K$ and $T_c = 305K$:

$$\eta_{max} = 1 - \frac{T_c}{T_h} = 1 - \frac{305K}{700K} = 1 - 0.4357 \approx 0.564 = 56.4\%$$

So the maximum theoretical efficiency between these temperatures is about 56.4%. The friend's claim of 95% therefore exceeds this ceiling by a substantial margin. By Carnot's theorem, no heat engine operating between 700K and 305K can have an efficiency above 56.4%, no matter how cleverly it is designed. The friend's claim is therefore impossible.

(b) The correct theoretical maximum efficiency is 56.4%, as computed in part (a). In physical terms, this means that for every 100J of heat absorbed by the engine from its hot reservoir at 700K, at most 56.4J can emerge as useful work, and at least 43.6J must be rejected to the cold reservoir at 305K as waste heat. Any engine that claims to deliver more than 56.4J of work per 100J of absorbed heat between these two reservoirs is claiming to violate the second law of thermodynamics, which is one of the most well-tested laws in all of physics. The friend's claim is therefore impossible regardless of the design details.

Making Sense of the Answer: *A 95% claim is not just slightly too high; it is too high by a factor that makes the device physically impossible, not merely difficult or inefficient. The second law is not an engineering challenge to be overcome with better machining or better fuel; it is a fundamental statement about what nature permits. A claim that exceeds the Carnot bound by even a percent or two would warrant scepticism; a claim that exceeds it by 38 percentage points warrants flat refusal. Real petrol-engine efficiencies in this temperature range run around 25% to 30%, well below the 56% Carnot ceiling, because of friction, incomplete combustion, and other practical inefficiencies. There is room for engineers to improve real engines toward the Carnot ceiling, but no room above it.*

Think Like a Physicist: *Whenever a problem presents an efficiency claim and two reservoir temperatures, the first move is to compute the Carnot bound. If the claim exceeds the bound, the device is impossible, full stop. If the claim falls below the bound, the device is in principle permitted, though it may still be difficult or impractical to build. The Carnot bound is the most economical possible filter for engine claims: one calculation, one comparison, and the verdict is in. Beware of impressive-sounding claims that do not respect the bound. Real engineering improvements work below the ceiling, never above it.*

HOT Example 38

A thermal power station in Tanzania burns heavy fuel oil. The combustion temperature in the boiler is 900K, and the cooling tower rejects waste heat to the surroundings at an ambient temperature of 310K.

(a) Calculate the maximum theoretical efficiency of the power station.
(b) The actual electrical efficiency achieved by the station is 38%. Calculate the fraction of fuel energy lost as waste heat, and explain in physical terms where this waste energy actually goes.
(c) Does the operation of this power station violate the first law of thermodynamics? Does it violate the second law? Explain each answer briefly.

Solution

(a) The maximum theoretical efficiency, by Carnot's theorem with T_h = 900K and T_c = 310K:

$$\eta_{max} = 1 - \frac{T_c}{T_h} = 1 - \frac{310K}{900K} = 1 - 0.3444 \approx 0.656,$$

or about 65.6%.

This is the absolute ceiling on the efficiency of any heat engine, including this power station, that operates between these two temperatures.

(b) The actual efficiency is 38%, meaning that 38% of the fuel's combustion energy emerges as useful electrical work. The remaining 62% is lost as waste heat:

$$\text{fraction lost} = 1 - 0.38 = 0.62 = 62\%.$$

Physically, this 62% of fuel energy is rejected to the atmosphere through the cooling tower, where it appears as warm water vapour rising from the tower and warm air around the station. None of this energy has been destroyed; the first law forbids that. It has been **degraded** from useful, high-temperature thermal energy (concentrated in the 900K boiler steam) to less useful, low-temperature thermal energy (dispersed in the 310K cooling water and surrounding air). The amount of energy is conserved, but its quality, or its capacity to do further work, has fallen.

(c) **Whether the power station violates the first law:** The first law states that energy is conserved in every process. The power station takes in 100J of fuel energy per arbitrary unit input, delivers 38J as electrical work, and rejects 62J as waste heat. The total output, 38J + 62J = 100J, equals the input. Energy is conserved exactly. The first law is not violated.

Whether the power station violates the second law: The second law, in the form of Carnot's theorem, sets the maximum possible efficiency between 900K and 310K at 65.6%. The actual efficiency, 38%, is below this ceiling. Therefore, the second law is not violated either. The power station operates within the bounds set by both laws, just well below the second-law ceiling because of practical inefficiencies (heat loss in pipes, friction in the turbine, incomplete combustion, and so on).

Making Sense of the Answer: *Three results emerge in tight agreement. The Carnot ceiling is 65.6%; the actual efficiency is 38%; the difference (27.6 percentage points) is the gap left by practical engineering, the room for improvement that engineers*

continually push against but cannot eliminate. The 62% of fuel energy not converted to electrical work is not destroyed; it is rejected as low-grade heat to the atmosphere. The cooling tower of every thermal power station in the country is the visible evidence of the second law: the plumes of warm vapour rising from the towers carry away the rejected heat that the second law insists must be discarded somewhere.

Think Like a Physicist: *The first law is concerned with **quantities** of energy; the second law is concerned with the **quality**, or the capacity, of that energy to do useful work. The same hundred joules can be high-quality if concentrated at a high temperature, or low-quality if dispersed at a low temperature. A power station's job is to convert as much high-quality input energy to useful work as possible before it inevitably becomes low-quality waste heat. The Carnot bound expresses the unavoidable price the second law extracts from this conversion, in proportion to the temperature ratio of cold to hot reservoirs. The closer the cold reservoir is to absolute zero, or the hotter the hot reservoir, the smaller this price. Practical engines aim for both: the highest combustion temperature the materials can stand, and the coldest cooling water available.*

This is where the chapter, and our introduction to thermodynamics, closes. The first law has been the constant companion of every problem: it has accounted for the conservation of energy in every process, from the simplest heating of a gas at constant volume to the most complicated multi-stage cycle. The second law, introduced briefly here, has supplied the directionality and the efficiency ceiling that the first law alone cannot deliver. Together, the two laws form the foundation of all classical thermodynamics, and although a full mathematical treatment of the second law lies in more advanced courses, the verbal statements and the Carnot bound that have been presented here are sufficient to answer every NECTA examination question on the second law that has been set in recent years. The physics of heat, work, and energy is a vast subject, and what has been built up in this chapter is only the entrance. But it is an entrance that opens directly onto every engine, refrigerator, and power station in the country, and on every cup of tea that cools toward room temperature.

The chapter's ten sections have each held one idea at a time. The miscellaneous worked examples that follow no longer have that courtesy. They will combine ideas across sections, and sometimes across chapters, in exactly the way examiners prefer.

MISCELLANEOUS WORKED EXAMPLES ON THERMODYNAMICS

Example 39

(a) Explain in molecular terms why a polyatomic ideal gas requires more heat than a monatomic ideal gas to undergo the same temperature rise at constant volume, even though both samples contain the same number of moles.

(b) 0.50 mol of an ideal monatomic gas absorbs 200J of heat. Calculate the temperature rise of the gas if (i) the process takes place at constant volume, (ii) the process takes place at constant pressure. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) A monatomic molecule, having no internal structure, can store kinetic energy only in three independent translational degrees of freedom. A polyatomic molecule, in addition to its three translational modes, can store kinetic energy in rotational modes about its molecular axes and in vibrational modes along its bonds. The equipartition theorem assigns an energy of $\frac{1}{2}kT$ to every accessible degree of freedom per molecule, so a polyatomic molecule must absorb energy not only into its translational motion but also into its rotational and vibrational modes before its temperature rises by the same amount. With more energy stores to fill per molecule, more heat is required for the same temperature rise.

(b)(i) At constant volume, all of the heat raises the internal energy of the gas, and for a monatomic gas $Q = nC_v\Delta T$ with $C_v = \frac{3}{2}R$.

$$\Delta T = \frac{Q}{nC_v} = \frac{200\text{J}}{0.50\text{mol} \times \frac{3}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1}} = 32.1\text{K}$$

(ii) At constant pressure, $Q = nC_p\Delta T$, with $C_p = C_v + R = \frac{3}{2}R + R = \frac{5}{2}R$ for a monatomic gas.

$$\Delta T = \frac{Q}{nC_p} = \frac{200\text{J}}{0.50\text{mol} \times \frac{5}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1}} = 19.2\text{K}$$

The temperature rise is approximately 32.1K at constant volume and 19.2K at constant pressure.

Example 40

(a) Distinguish, by reference to the first law of thermodynamics, between the work done by an ideal gas during an isothermal expansion and the work done during an adiabatic expansion when both processes start from the same initial state and produce the same volume change. Account for which of the two processes delivers the larger amount of work, and explain why.

(b) 2.0 mol of an ideal gas ($\gamma = 1.4$) initially at 350K occupies a volume of 0.020m^3 and expands to a volume of 0.040m^3. Calculate the work done by the gas if (i) the expansion is isothermal at 350K, (ii) the expansion is adiabatic. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) In an isothermal expansion the temperature is held constant, so the internal energy of an ideal gas does not change, and the first law $\Delta U = Q - W$ reduces to $Q = W$; every joule of heat that flows in is delivered as work. In an adiabatic expansion no heat enters the gas, so $\Delta U = -W$, and the work is supplied entirely from the gas's own internal energy. As the gas does work adiabatically its internal energy and therefore its temperature fall, the pressure drops faster than the corresponding isotherm, and the area under the adiabatic PV-curve is smaller than the area under the isothermal curve over the same volume range. The isothermal expansion therefore delivers more work than the adiabatic expansion between the same initial state and the same final volume.

(b)(i) For the isothermal expansion at constant $T = 350K$, the work done by the gas is $W = nRT\ln V_2/V_1$.

$$W = nRT\ln\frac{V_2}{V_1} = 2.0mol \times 8.314Jmol^{-1}K^{-1} \times 350K \times \ln\frac{0.040m^3}{0.020m^3} = 4034J$$

(ii) For the adiabatic expansion, $TV^{\gamma-1} = constant$, gives the final temperature.

$$T_2 = T_1\left(\frac{V_1}{V_2}\right)^{\gamma-1} = 350K \times \left(\frac{0.020}{0.040}\right)^{0.4} = 265.3K$$

Since $Q = 0$, the work done by the gas equals the loss of internal energy:

$$W = nC_v(T_1 - T_2) \text{ with } C_v = R/(\gamma - 1) = \frac{5}{2}R$$

$$W = nC_v(T_1 - T_2) = 2.0mol \times \frac{5}{2} \times 8.314Jmol^{-1}K^{-1} \times (350K - 265.3K) = 3520J$$

The work done is approximately 4034J in the isothermal expansion and 3520J in the adiabatic expansion, and the isothermal value is the larger of the two as predicted.

Example 41

(a) When a small mass of water at 100°C is fully vapourised at atmospheric pressure, only a small fraction of the heat supplied appears as work done by the steam against the atmosphere, while the larger fraction is stored as an increase in internal energy of the water-steam system. Account for this partition of energy in physical terms.

(b) 5.0g of water at 100°C is fully vapourised at atmospheric pressure (1.013×10^5Pa). Calculate the work done by the steam as it forms, the increase in internal energy of the water-steam system, and the total heat supplied. Take the specific latent heat of vapourisation of water as $L_v = 2.26 \times 10^6 Jkg^{-1}$, the density of liquid water as $\rho_l = 1000kgm^{-3}$, and the density of steam at 100°C and atmospheric pressure as $\rho_s = 0.598kgm^{-3}$.

Solution

(a) In the liquid state the water molecules sit close together, held by strong intermolecular (attractive) forces, so liquid water has a small specific volume. In the vapour state the molecules are far apart and the gas occupies a large specific volume. Vapourising water therefore requires energy on two distinct accounts. First, the molecules must be pulled apart against their mutual attraction, a process that increases the internal energy of the system without producing any external mechanical effect. Second, the new vapour pushes back the surrounding atmosphere as it expands into the space it now needs, doing mechanical work on the surroundings. Because the intermolecular attractions in liquid water are strong relative to the modest pressure of the atmosphere, the energy needed to overcome these attractions is much larger than the work needed to displace the atmosphere; the partition is therefore heavily weighted towards internal energy.

(b) Volume of the liquid: $V_l = m/\rho_l = 5.0 \times 10^{-3}kg / 1000kgm^{-3} = 5.0 \times 10^{-6}m^3$. Volume of the steam at 1.013×10^5Pa and 100°C: $V_s = m/\rho_s = 5.0 \times 10^{-3}kg / 0.598kgm^{-3} = 8.36 \times 10^{-3}m^3$. The work done by the steam against the atmosphere as it forms is $W = P\Delta V$ at constant atmospheric pressure.

$$W = P(V_s - V_l) = 1.013 \times 10^5 Pa \times (8.36 \times 10^{-3}m^3 - 5.0 \times 10^{-6}m^3) = 846J$$

The total heat supplied is $Q = mL_v$, by the definition of the specific latent heat of vapourisation at constant pressure.

$$Q = mL_v = 5.0 \times 10^{-3}kg \times 2.26 \times 10^6 Jkg^{-1} = 11300J$$

From the first law of thermodynamics, the increase in internal energy is $\Delta U = Q - W$.

$$\Delta U = Q - W = 11300J - 846J = 10454J$$

The work done by the steam against the atmosphere is approximately 846J, the increase in internal energy is approximately 10454J, and the total heat supplied is approximately 11300J. Of the heat supplied, only about 7.5% emerges as external work; the remaining 92.5% is stored as internal energy.

Example 42

(a) Justify, by reference to the equipartition of energy, why a polyatomic ideal gas with seven accessible degrees of freedom (three translational, three rotational, and one vibrational) has a value of γ closer to unity than that of a monatomic ideal gas. Determine the value of γ for such a polyatomic gas.

(b) 3.0 mol of the polyatomic gas described above is contained at a temperature of 400K and undergoes an isothermal expansion in which its volume doubles. Calculate the work done by the gas during this expansion. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) By the equipartition theorem each accessible degree of freedom stores an average energy of $\frac{1}{2}\text{kT}$ per molecule, so the molar internal energy of an ideal gas with f accessible degrees of freedom is $U = \frac{f}{2}RT$, the molar heat capacity at constant volume is $C_v = \frac{f}{2}R$, and the molar heat capacity at constant pressure is $C_p = C_v + R = \frac{f+2}{2}R$. The ratio of principal heat capacities is therefore $\gamma = C_p/C_v = (f+2)/f$. For $f = 3$ (monatomic) this gives $\gamma = 5/3 \approx 1.67$. For $f = 7$ (polyatomic with three translational, three rotational, and one vibrational degree of freedom) this gives $\gamma = 9/7 \approx 1.29$, a value closer to unity. As f rises, the additional energy stores dilute the rise in temperature per joule of heat absorbed, the difference between C_p and C_v becomes a smaller fraction of either, and γ approaches one.

$$\gamma = \frac{f+2}{f} = \frac{7+2}{7} = \frac{9}{7} \approx 1.29$$

(b) For an isothermal expansion the work done by the gas is $W = nRT\ln V_2/V_1$, independent of γ.

$$W = nRT\ln\frac{V_2}{V_1} = 3.0\text{mol} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 400\text{K} \times \ln 2 = 6915\text{J}$$

The value of γ for the polyatomic gas is $9/7 \approx 1.29$, and the work done in the isothermal expansion is approximately 6915J.

Example 43

(a) Identify the three distinct contributions to the total heat input when a fixed mass of liquid water in an open container is heated from below room temperature, brought to its boiling point, partially vapourised, and the new steam is allowed to expand against atmospheric pressure as it forms.

(b) 0.50kg of water in an open container is heated from 25°C and partially boiled away at atmospheric pressure (1.013×10^5Pa) until 0.30kg of liquid water remains. Calculate: (i) the total heat supplied to the system, (ii) increase in internal energy; given the specific heat capacity of water $c = 4200\text{Jkg}^{-1}\text{K}^{-1}$, the specific latent heat of vaporisation $L_v = 2.26 \times 10^6\text{Jkg}^{-1}$, and the density of steam at 100°C and 1 atmosphere $\rho_s = 0.598\text{kgm}^{-3}$. Neglect the heat capacity of the container.

Solution

(a) The **first** contribution is the *sensible heat* needed to raise the temperature of the whole liquid mass from 25°C to its boiling point of 100°C, given by mcΔT. The **second** contribution is the energy that increases the internal energy of the vapourised portion, *breaking the intermolecular attractions* in the liquid as molecules escape into the vapour phase. The **third** contribution is the **mechanical work** done by the new steam as it forms and pushes back the surrounding atmosphere by an amount equal to the volume of the new vapour. The standard latent heat of vapourisation, quoted at constant atmospheric pressure, packages the second and third contributions into a single coefficient, so the total heat supplied is conveniently written as the sum of the sensible heat and the latent heat term.

(b) Sensible heat to raise the liquid from 25°C to 100°C, with $m = 0.50\text{kg}$, $c = 4200\text{Jkg}^{-1}\text{K}^{-1}$ and $\Delta T = 75\text{K}$:

$$Q_1 = mc\Delta T = 0.50\text{kg} \times 4200\text{Jkg}^{-1}\text{K}^{-1} \times 75\text{K} = 157500\text{J}$$

Mass vapourised: $\Delta m = 0.50\text{kg} - 0.30\text{kg} = 0.20\text{kg}$. Latent heat associated with the vaporisation:

$$Q_2 = \Delta m \times L_v = 0.20\text{kg} \times 2.26\times10^6\text{Jkg}^{-1} = 4.52\times10^5\text{J}$$

Within Q_2, the work done by the new steam against the atmosphere is $W = P\Delta V$, where ΔV is the volume occupied by the steam after it forms, $V = \Delta m/\rho_s$.

$$W = P \times \frac{\Delta m}{\rho_s} = 1.013\times10^5\text{Pa} \times \frac{0.20\text{kg}}{0.598\text{kgm}^{-3}} = 33880\text{J}$$

The remaining portion of Q_2, namely $Q_2 - W = 4.52\times10^5\text{J} - 3.39\times10^4\text{J} = 4.18\times10^5\text{J}$, is stored as the **internal energy increase** associated with breaking the intermolecular attractions during vapourisation. The **total heat supplied** is the sum of the sensible heat and the latent heat term:

$$Q_{\text{total}} = Q_1 + Q_2 = 1.575\times10^5\text{J} + 4.52\times10^5\text{J} = 6.10\times10^5\text{J}$$

The **total heat** supplied to the system is approximately $\mathbf{6.10\times10^5 J}$, of which, 1.575×10^5J (≈26%) raises the temperature of the liquid, $\mathbf{4.18\times10^5 J}$ (≈68%) increases the **internal energy** of the new vapour, and 3.39×10^4J (≈6%) is delivered as work against the atmosphere.

Example 44

(a) Explain why, when heat is supplied to a gas in a sealed rigid container of metal walls in contact with a hotter reservoir, the rate at which the internal energy of the gas rises depends on both the thermal conductivity of the container walls and the cross-sectional area available for conduction.

(b) A sealed rigid steel container of internal volume 0.020m^3 contains 0.80 mol of monatomic ideal gas, initially at 27°C. One end of the container, of cross-sectional area 0.010m^2 and wall thickness 5.0mm, is in steady contact with a steam reservoir at 100°C, while the rest of the container surface is perfectly insulated. The thermal conductivity of steel is $k = 50\text{Wm}^{-1}\text{K}^{-1}$. Calculate the initial rate at which the temperature of the gas rises with time. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) Heat reaches the gas only through the conducting wall, and Fourier's law fixes the rate at which energy crosses that wall as $dQ/dt = kA\Delta T/L$, where k is the thermal conductivity of the wall material, A is the cross-sectional area through which heat flows, ΔT is the temperature difference across the wall, and L is the wall thickness. Increasing k means each square metre of wall passes heat more readily; increasing A increases the total area available for heat passage; and both raise the rate at which energy enters the gas. Because the container is rigid the gas does no work, so by the first law every joule arriving in the gas is delivered to its internal energy, and the rate of temperature rise dT/dt is proportional to dQ/dt. The two material parameters of the wall therefore appear directly in the rate at which the gas heats up.

(b) Temperature difference across the wall: $\Delta T = 373\text{K} - 300\text{K} = 73\text{K}$. Initial rate of heat conduction into the gas:

$$\frac{dQ}{dt} = \frac{kA\Delta T}{L} = \frac{50\text{Wm}^{-1}\text{K}^{-1} \times 0.010\text{m}^2 \times 73\text{K}}{5.0 \times 10^{-3}\text{m}} = 7300\text{W}$$

At constant volume the heat raises only the internal energy: $dQ = nC_v dT$, with $C_v = \frac{3}{2}R$ for the monatomic gas, hence:

$$\frac{dT}{dt} = \frac{1}{nC_v} \times \frac{dQ}{dt} = \frac{7300\text{W}}{0.80\text{mol} \times \frac{3}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1}} = 732\text{Ks}^{-1}$$

The initial rate at which the temperature of the gas rises is approximately 732Ks^{-1}.
The rate is large because the small quantity of monatomic gas has a small heat capacity while the steel wall is thin and conducts well.

Example 45

(a) For an ideal gas undergoing a quasi-static adiabatic process, derive the relation $\frac{dT}{T} + (\gamma - 1)\frac{dV}{V} = 0$.

(b) 2.0 mol of an ideal gas ($\gamma = 1.4$) is initially at 320K and 0.050m^3. Using the relation derived in part (a), calculate the temperature of the gas when its volume has decreased to 0.020m^3 during a quasi-static adiabatic compression.

Solution

(a) In an adiabatic process $Q = 0$, so the first law $dU = dQ - dW$ for an infinitesimal step reduces to $dU = -dW$. For a quasi-static process the work done by the gas in expanding by dV against its own pressure is $dW = PdV$, and the internal energy of an ideal gas depends only on temperature, with $dU = nC_v dT$. Combining these gives $nC_v dT = -PdV$. Substituting $P = nRT/V$ from the ideal gas equation and dividing through by nT yields:

$$\frac{dT}{T} = -\frac{R}{C_v} \times \frac{dV}{V}$$

The combination R/C_v equals $(C_p - C_v)/C_v = \gamma - 1$, so $dT/T = -(\gamma - 1)dV/V$. Rearranging gives the required result:

$$\frac{dT}{T} + (\gamma - 1)\frac{dV}{V} = 0$$

(b) Integrating the differential relation between the initial state (T_1, V_1) and the final state (T_2, V_2):

$$\int_{T_1}^{T_2} \frac{dT}{T} = -(\gamma - 1)\int_{V_1}^{V_2} \frac{dV}{V}$$

$$\ln\frac{T_2}{T_1} = -(\gamma - 1)\ln\frac{V_2}{V_1}$$

$$\ln\frac{T_2}{T_1} = \ln\left(\frac{V_2}{V_1}\right)^{-(\gamma-1)} = \ln\left(\frac{V_1}{V_2}\right)^{(\gamma-1)}$$

$$\frac{T_2}{T_1} = \left(\frac{V_1}{V_2}\right)^{(\gamma-1)}$$

Solving for T_2 with $T_1 = 320\text{K}$, $V_1 = 0.050\text{m}^3$, $V_2 = 0.020\text{m}^3$ and $\gamma = 1.4$:

$$T_2 = T_1\left(\frac{V_1}{V_2}\right)^{\gamma-1} = 320\text{K} \times \left(\frac{0.050}{0.020}\right)^{0.4} = 462\text{K}$$

The temperature of the gas at the end of the adiabatic compression is approximately 462K.

Example 46

(a) A refrigerator powered by an electrical supply is placed in a perfectly sealed and perfectly insulated room and is left running with its door open. Predict, with reason from the first law of thermodynamics, what happens to the average temperature of the room over time.

(b) The refrigerator described in part (a) is rated at an electrical input of 150W. Estimate the rate at which the temperature of the air in the room rises if the room has internal dimensions $5.0\text{m} \times 4.0\text{m} \times 3.0\text{m}$ and contains air of density $\rho = 1.2\text{kgm}^{-3}$ and specific heat capacity at constant volume $c_v = 720\text{Jkg}^{-1}\text{K}^{-1}$. Neglect the heat capacity of the refrigerator and of the contents of the room.

Solution

(a) Treating the room and the refrigerator together as a single closed and insulated system, the only energy crossing the boundary of the system is the electrical energy supplied through the cable. The first law applied to the whole system therefore gives $\Delta U_{\text{system}} = W_{\text{electrical}}$, so the entire electrical input goes into the internal energy of the air in the room. The refrigerator does not destroy any energy by removing heat from its interior; it only relocates heat, taking it from inside the cabinet (which the open door now exposes to the room) and delivering it, together with the work it consumes, back into the same room. The average temperature of the room therefore rises steadily as the compressor runs.

(b) Volume of the room: $V = 5.0\text{m} \times 4.0\text{m} \times 3.0\text{m} = 60\text{m}^3$.
Mass of air: $m = \rho V = 1.2\text{kgm}^{-3} \times 60\text{m}^3 = 72\text{kg}$.
The rate at which the air absorbs heat at constant volume equals the electrical power input:

$$mc_v \frac{dT}{dt} = P_{\text{elec}}$$

$$\frac{dT}{dt} = \frac{P_{\text{elec}}}{mc_v} = \frac{150\text{W}}{72\text{kg} \times 720\text{Jkg}^{-1}\text{K}^{-1}} = 2.89 \times 10^{-3}\text{Ks}^{-1}$$

The temperature of the room rises at approximately $2.89\times10^{-3}\text{Ks}^{-1}$.

This is about 0.17K every minute or 10K every hour, sufficient to be uncomfortable within an hour and dangerous within a working day.

Example 47

(a) Show, with reasons, that the net heat absorbed by an ideal gas during one complete cycle of any closed thermodynamic process is necessarily equal to the net work done by the gas during that cycle.

(b) 1.5 mol of an ideal gas ($\gamma = 1.4$) executes the following cycle. Starting at state A with $P = 1.0 \times 10^5\text{Pa}$, $V = 0.040\text{m}^3$, $T = 320\text{K}$, the gas undergoes (i) an isobaric expansion to $V = 0.080\text{m}^3$, (ii) an isochoric cooling until the pressure has fallen to $0.50 \times 10^5\text{Pa}$, (iii) an isothermal compression back to state A. Calculate the net work done by the gas during the complete cycle. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) For any cycle the gas returns to its starting state, so every state variable, including the internal energy, returns to its initial value and $\Delta U_{\text{cycle}} = 0$. The first law for the complete cycle therefore reduces to $Q_{\text{net}} = W_{\text{net}}$: every joule of net heat absorbed during the cycle is delivered as net work, and the two are necessarily equal whatever the shape of the cycle, whatever the working substance, and whatever the sequence of processes. The equality is a direct consequence of internal energy being a state function.

(b) The three legs of the cycle. Leg (i) isobaric A → B at $P = 1.0 \times 10^5\text{Pa}$, expansion from $V = 0.040\text{m}^3$ to $V = 0.080\text{m}^3$:

$$W_1 = P\Delta V = 1.0 \times 10^5\text{Pa} \times (0.080\text{m}^3 - 0.040\text{m}^3) = 4000\text{J}$$

At B the temperature has risen to $T_B = T_A \times V_B/V_A = 320\text{K} \times 2 = 640\text{K}$.

Leg (ii) isochoric B → C at $V = 0.080\text{m}^3$, cooling until $P = 0.50 \times 10^5\text{Pa}$: no volume change, so $W_2 = 0$, and $T_C = T_B \times P_C/P_B = 640\text{K} \times 0.50 = 320\text{K}$.

The gas at C is therefore at the same temperature as at A, so leg (iii) is the isothermal compression at 320K from $V_C = 0.080\text{m}^3$ to $V_A = 0.040\text{m}^3$:

$$W_3 = nRT\ln\frac{V_A}{V_C} = 1.5\text{mol} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 320\text{K} \times \ln\frac{0.040}{0.080} = -2766\text{J}$$

The negative sign indicates that work is done on the gas during the compression. The net work for the complete cycle is the sum of the three:

$$W_{\text{net}} = W_1 + W_2 + W_3 = 4000\text{J} + 0 + (-2766\text{J}) = 1234\text{J}$$

The net work done by the gas during the complete cycle is approximately 1234J (and by the cycle closure rule this is also the net heat absorbed during the cycle).

Example 48

(a) In industrial gas-compression machinery a gas is often compressed in two adiabatic stages, with isobaric cooling between the stages, rather than in a single adiabatic stage. Identify the physical reason engineers prefer the two-stage arrangement, and account for it in terms of the work that must be supplied to the compressor.

(b) 0.50 mol of an ideal gas ($\gamma = 1.4$) initially at 1.0×10^5Pa and 300K is compressed in two stages. Stage 1: adiabatic compression until the volume halves. Stage 2: isobaric cooling at the new pressure until the temperature returns to 300K. Calculate the total work done on the gas during the two stages combined, and the heat rejected during the isobaric cooling stage. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) A single-stage adiabatic compression heats the gas substantially as it does work on it, so the gas is then at a high temperature and high pressure. The work needed to push it the rest of the way to a target final pressure is large because the pressure is already large. Cooling the gas back to its starting temperature between stages, at the elevated intermediate pressure, removes the temperature rise without changing the pressure, so the second adiabatic stage starts from a colder gas at the same pressure and ends at the target final pressure with a smaller total temperature rise and a smaller swept volume. The two areas under the two adiabats on the PV diagram together cover less area than the single adiabat to the same final pressure, so the integrated work supplied to the compressor is reduced. The intercooling step rejects to the surroundings the heat that the first-stage compression deposited in the gas, and that heat would otherwise have to be paid for as additional compressor work in the single-stage arrangement.

(b) Initial volume of the gas: $V_1 = \frac{nRT_1}{P_1} = \frac{0.50\text{ mol} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 300\text{K}}{1.0\times10^5\text{Pa}} = 1.247 \times 10^{-2}\text{m}^3$.

Stage 1, adiabatic compression to $V_2 = \frac{V_1}{2} = 6.236 \times 10^{-3}\text{m}^3$:

$$T_2 = T_1\left(\frac{V_1}{V_2}\right)^{\gamma-1} = 300\text{K} \times 2^{0.4} = 396\text{K}$$

$$P_2 = P_1\left(\frac{V_1}{V_2}\right)^{\gamma} = 1.0 \times 10^5\text{Pa} \times 2^{1.4} = 2.64 \times 10^5\text{Pa}$$

Work done on the gas in the adiabatic stage equals the increase in internal energy, with $C_v = \frac{R}{\gamma-1} = \frac{5}{2}R$:

$$W_{1,on} = nC_v(T_2 - T_1) = 0.50\text{mol} \times \frac{5}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times (396\text{K} - 300\text{K}) = 996\text{J}$$

Stage 2 is isobaric cooling at P_2 from $T_2 = 396K$ back to 300K.

The volume falls from $V_2 = 6.236 \times 10^{-3}\text{m}^3$ to $V_3 = V_2 \times \frac{T_3}{T_2} = 6.236 \times 10^{-3} \times \frac{300}{396} = 4.726 \times 10^{-3}\text{m}^3$.

Work done on the gas:

$$W_{2,on} = -P_2(V_3 - V_2) = -2.64 \times 10^5\text{Pa} \times (4.726 - 6.236) \times 10^{-3}\text{m}^3 = 398\text{J}$$

Total work done on the gas in the two stages combined:

$$W_{on,total} = W_{1,on} + W_{2,on} = 996\text{J} + 398\text{J} = 1394\text{J}$$

Heat rejected during the isobaric cooling, $Q_{rejected} = nC_p(T_2 - T_3)$ with $C_p = \frac{7}{2}R$:

$$Q_{rejected} = nC_p(T_2 - T_3) = 0.50\text{mol} \times \frac{7}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times (396\text{K} - 300\text{K}) = 1395\text{J}$$

The total work done on the gas during the two stages is approximately 1394J and the heat rejected during the isobaric cooling stage is approximately 1395J.

Example 49

(a) A piston of mass m and cross-sectional area A rests on top of a fixed mass of gas in a vertical cylinder open to the atmosphere above and is free to slide vertically. By a force-balance argument on the piston, derive the pressure of the gas inside the cylinder in terms of the atmospheric pressure, the piston's weight and the area, and explain why this gas pressure is independent of the piston's vertical position, provided the piston is in mechanical equilibrium.

(b) A vertical cylinder of internal cross-sectional area $A = 0.020\text{m}^2$ is closed below and contains 0.10 mol of an ideal gas. A piston of mass $m = 5.0\text{kg}$ rests on top of the gas, free to slide vertically. The atmospheric pressure is $P_{atm} = 1.013 \times 10^5$Pa and gravity is $g = 9.81\text{ms}^{-2}$. The gas is heated and the piston rises by $\Delta h = 0.10\text{m}$. Calculate the work done by the gas during this process, identifying separately the work done against the atmosphere and the work done in raising the piston. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) A free piston in mechanical equilibrium has zero net vertical force on it. The downward forces are its weight mg and the atmospheric pressure P_{atm} acting over its top surface, while the upward force is the gas pressure acting over its bottom surface. Equating these on a piston of cross-sectional area A gives:

$$P_{gas}A = P_{atm}A + mg, \text{hence } P_{gas} = P_{atm} + \frac{mg}{A}$$

The result depends only on the piston's weight, the atmospheric pressure and the area, none of which involves the piston's vertical position; the gas pressure is therefore the same whether the piston sits low in the cylinder or high in it.

(b) Pressure of the gas in the cylinder:

$$P_{gas} = P_{atm} + \frac{mg}{A} = 1.013 \times 10^5 \text{Pa} + \frac{5.0\text{kg} \times 9.81\text{ms}^{-2}}{0.020\text{m}^2} = 1.038 \times 10^5 \text{Pa}$$

Volume swept by the piston: $\Delta V = A\Delta h = 0.020\text{m}^2 \times 0.10\text{m} = 2.0 \times 10^{-3}\text{m}^3$.

Total work done by the gas at constant gas pressure:

$$W_{total} = P_{gas}\Delta V = 1.038 \times 10^5 \text{Pa} \times 2.0 \times 10^{-3}\text{m}^3 = 207.5\text{J}$$

Work done by the gas against the atmosphere:

$$W_{atm} = P_{atm}\Delta V = 1.013 \times 10^5 \text{Pa} \times 2.0 \times 10^{-3}\text{m}^3 = 202.6\text{J}$$

Work done by the gas in raising the piston (equal to the gain in gravitational potential energy of the piston):

$$W_{piston} = mg\Delta h = 5.0\text{kg} \times 9.81\text{ms}^{-2} \times 0.10\text{m} = 4.91\text{J}$$

The total work done by the gas is approximately 207.5J, of which approximately 202.6J pushes back the atmosphere and approximately 4.91J raises the piston, in the expected ratio $P_{atm}: mg/A$.

Example 50

(a) The speed of sound in an ideal gas is given by $v = \sqrt{\gamma RT/M}$, where γ is the ratio of principal heat capacities. Account for the appearance of the factor γ, rather than the simpler factor 1, in this formula by reference to the time-scale on which compressions and rarefactions in a sound wave act on the gas.

(b) Calculate the speed of sound in dry air at 27°C, treating air as a diatomic ideal gas with $\gamma = 1.4$ and effective molar mass $M = 29 \times 10^{-3}\text{kgmol}^{-1}$. Compare the result with the root mean square speed of air molecules at the same temperature, and comment on the ratio. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) A sound wave is a sequence of compressions and rarefactions that pass through a gas at audio frequencies, sweeping any small region of gas through one pressure cycle in a time of order one millisecond. Heat conduction across the wavelength of an audible sound wave acts on a much longer time-scale, so essentially no heat flows in or out of any particular region during one cycle; the local compressions and rarefactions are therefore adiabatic rather than isothermal. The bulk modulus that controls the wave speed is consequently the adiabatic bulk modulus γP, not the isothermal value P, and in the formula $v = \sqrt{B/\rho}$ the factor γ appears. If the propagation were isothermal one would have $v = \sqrt{RT/M}$, the simpler formula with γ replaced by unity, but acoustic measurements unambiguously favour the larger adiabatic value, vindicating the kinetic-theory picture.

(b) Speed of sound at 300K:

$$v_{sound} = \sqrt{\frac{\gamma RT}{M}} = \sqrt{\frac{1.4 \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 300\text{K}}{29 \times 10^{-3}\text{kgmol}^{-1}}} = 347\text{ms}^{-1}$$

Root mean square speed at the same temperature:

$$v_{rms} = \sqrt{\frac{3RT}{M}} = \sqrt{\frac{3 \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 300\text{K}}{29 \times 10^{-3}\text{kgmol}^{-1}}} = 508\text{ms}^{-1}$$

Ratio of the two speeds:

$$\frac{v_{sound}}{v_{rms}} = \sqrt{\frac{\gamma}{3}} = \sqrt{\frac{1.4}{3}} = 0.683$$

The speed of sound in dry air at 27°C is approximately 347ms^{-1}, against an rms molecular speed of approximately 508ms^{-1}; the sound speed is about 68% of the rms molecular speed. The two are of the same order of magnitude because both are set by the same molecular kinetic energy scale, but they are not equal because the sound wave samples a typical molecular speed weighted by its ability to transfer momentum forward, not the bare rms speed.

Example 51

(a) In a vigorous downstroke of a bicycle pump, the air trapped in the barrel is compressed rapidly. Justify why this compression can be modelled as adiabatic, and account for the physical reason the metal barrel of the pump becomes warm to the touch after a few strokes.

(b) A cyclist pumps her bicycle tyre. Just before one downstroke, the air in the pump barrel occupies a volume of $4.0 \times 10^{-4}\text{m}^3$ at 27°C and atmospheric pressure ($1.013 \times 10^5\text{Pa}$). After the downstroke, the air has been compressed to a volume of $1.0 \times 10^{-4}\text{m}^3$, and a thermometer in good thermal contact with the air reads 187°C immediately after. Treating

the compression as adiabatic, calculate the value of γ for the air, and compare it with the value expected for a diatomic ideal gas.

Solution

(a) The downstroke takes a small fraction of a second, while heat conduction through the air to the metal barrel and out to the surroundings happens on a much longer time-scale, so during the stroke very little heat leaves the air. The compression is therefore approximately adiabatic, and the work done on the gas by the cyclist appears as an increase in internal energy, raising the temperature. After the stroke the now-hot air slowly transfers heat to the cooler metal barrel by conduction; since this transfer is repeated stroke after stroke, the barrel accumulates the deposited heat and becomes noticeably warm to the touch.

(b) For an adiabatic process $T_1 V_1^{\gamma-1} = T_2 V_2^{\gamma-1}$. Taking logarithms and rearranging gives:

$$\gamma - 1 = \frac{\log T_2/T_1}{\log V_1/V_2}$$

With $T_1 = 300\text{K}$, $T_2 = 460\text{K}$, $V_1 = 4.0 \times 10^{-4}\text{m}^3$ and $V_2 = 1.0 \times 10^{-4}\text{m}^3$:

$$\gamma - 1 = \frac{\log 460/300}{\log 4.0/1.0} = 0.308$$

$$\gamma = 1 + 0.308 = 1.31$$

The deduced value of γ is approximately 1.31, lower than the textbook value of 1.40 for an ideal diatomic gas. The discrepancy is consistent with the compression not being perfectly adiabatic; a small leak of heat from the gas to the metal barrel during the stroke pulls the apparent γ down towards the isothermal limit of unity.

Example 52

(a) An ideal gas occupies one half of a thermally insulated container of total volume $1.0 \times 10^{-2}\text{m}^3$. The other half is initially evacuated. The partition between the two halves is suddenly removed, and the gas freely fills the whole container. Show, using the first law of thermodynamics, that the temperature of the gas does not change during this process, even though its volume doubles. Account for why the equation $PV^{\gamma} = \text{constant}$ does **not** apply here.

(b) 0.20 mol of an ideal gas ($\gamma = 1.4$) is initially confined at 27°C and $1.0 \times 10^5\text{Pa}$ to one half of the $1.0 \times 10^{-2}\text{m}^3$ thermally insulated container described in part (a), with the other half evacuated. The partition is suddenly removed. Calculate (i) the final temperature of the gas, (ii) the final pressure of the gas, (iii) the change in internal energy of the gas.

Solution

(a) The container is thermally insulated, so no heat enters or leaves the gas, $Q = 0$. The other half of the container is evacuated, so the gas does no work on anything as it expands into the vacuum, $W = 0$. The first law $\Delta U = Q - W$ therefore gives $\Delta U = 0$. For an ideal gas the internal energy depends only on temperature, so $\Delta U = 0$ implies $\Delta T = 0$ and the temperature is unchanged at 300K. The relation $PV^{\gamma} = \text{constant}$ is derived under the assumption of a quasi-static adiabatic process, in which the gas is in equilibrium at every step and does work PdV at every step. Joule's free expansion is adiabatic but not quasi-static and does no work, so the derivation does not apply, and the gas does not follow an adiabat on the PV diagram; it simply jumps from its initial state to its final state along a horizontal isotherm.

(b)(i) Final temperature, by the argument of part (a):

$$T_f = T_i = 300\text{K}$$

(ii) For an ideal gas at constant temperature, PV is constant, so $P_f V_f = P_i V_i$, with $V_f = 2V_i$:

$$P_f = P_i \times \frac{V_i}{V_f} = 1.0 \times 10^5\text{Pa} \times \frac{V_i}{2V_f} = 5.0 \times 10^4\text{Pa}$$

(iii) Change in internal energy:

$$\Delta U = nC_v\Delta T = nC_v \times 0 = 0$$

The final temperature of the gas is 300K, the final pressure is 5.0×10^4Pa, and the change in internal energy is zero, all in agreement with the qualitative argument from the first law.

Example 53

(a) A closed cycle on a PV diagram consists of two curves of one type joined by two curves of a different type, and the cycle is traversed clockwise. Identify what the two types of curve must be in order for the cycle to operate as a Carnot engine, and justify your identification by reference to the requirement that all heat absorbed and rejected by a Carnot engine occurs at exactly two temperatures.

(b) A Carnot engine operates between a hot reservoir at 500K and a cold reservoir at 300K. During one complete cycle, the engine absorbs 1500J of heat from the hot reservoir. Calculate (i) the efficiency of the engine, (ii) the work output per cycle, (iii) the heat rejected to the cold reservoir per cycle.

Solution

(a) A Carnot engine, by Carnot's construction, exchanges heat with the surroundings only at exactly two temperatures, the hot-reservoir temperature T_h and the cold-reservoir temperature T_c. Heat absorption from the hot reservoir at constant T_h occurs along an isothermal expansion at T_h; heat rejection to the cold reservoir at constant T_c occurs along an isothermal compression at T_c. The two isotherms are joined by processes during which no heat is exchanged with anything, that is, by adiabatic legs: one expansion taking the gas from T_h down to T_c and one compression returning it from T_c up to T_h. The Carnot cycle therefore consists of two isothermal curves joined by two adiabatic curves, and these are the only curves consistent with heat exchange occurring at exactly two temperatures.

(b)(i) Efficiency of the Carnot engine:

$$\eta = 1 - \frac{T_c}{T_h} = 1 - \frac{300\text{K}}{500\text{K}} = 0.40$$

(ii) Work output per cycle:

$$W = \eta \times Q_h = 0.40 \times 1500\text{J} = 600\text{J}$$

(iii) Heat rejected per cycle, by the first law for the cycle:

$$Q_c = Q_h - W = 1500\text{J} - 600\text{J} = 900\text{J}$$

The efficiency of the Carnot engine is 40%, the work output per cycle is 600J, and the heat rejected to the cold reservoir per cycle is 900J.

Example 54

(a) An ideal gas undergoes a process in which its temperature varies with volume according to $T = T_0 + \alpha V^2$, where T_0 and α are positive constants and V is the volume of the gas. Determine, in terms of T_0, α and the universal gas constant R, the volume at which the pressure of one mole of the gas reaches its minimum value, and derive an expression for the corresponding minimum pressure.

(b) For one mole of helium gas with $T_0 = 250\text{K}$ and $\alpha = 1.2 \times 10^6\text{Km}^{-6}$, calculate (i) the volume at which the pressure of the gas is a minimum, (ii) the minimum pressure attained, (iii) the temperature of the gas at the minimum pressure. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Solution

(a) For one mole the ideal gas equation gives $PV = RT$, hence $P = RT/V$. Substituting $T = T_0 + \alpha V^2$ yields P as an explicit function of V:

$$P(V) = \frac{R(T_0 + \alpha V^2)}{V} = \frac{RT_0}{V} + R\alpha V$$

The first term decreases with V while the second increases with V, so P(V) has a single minimum. Setting $dP/dV = 0$ to locate it:

$$\frac{dP}{dV} = -\frac{RT_0}{V^2} + R\alpha = 0$$

$$\mathbf{V = \sqrt{\frac{T_0}{\alpha}}}$$

Substituting V back into P(V) gives the minimum pressure:

$$\mathbf{P_{min}} = \frac{RT_0}{V} + R\alpha V = R\sqrt{\alpha T_0} + R\sqrt{\alpha T_0} = \mathbf{2R\sqrt{\alpha T_0}}$$

(b)(i) With $T_0 = 250\text{K}$ and $\alpha = 1.2 \times 10^6\text{Km}^{-6}$:

$$V_{min} = \sqrt{\frac{250\text{K}}{1.2 \times 10^6\text{Km}^{-6}}} = \sqrt{2.083 \times 10^{-4}\text{m}^6} = 1.44 \times 10^{-2}\text{m}^3$$

(ii) Minimum pressure:

$$P_{min} = 2R\sqrt{\alpha T_0} = 2 \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times \sqrt{1.2 \times 10^6\text{Km}^{-6} \times 250\text{K}} = 2.88 \times 10^5\text{Pa}$$

(iii) At P_{min} the temperature is $T = T_0 + \alpha V^2 = T_0 + \alpha(T_0/\alpha) = 2T_0$

$$2T_0 = 2 \times 250\text{K} = 500\text{K}$$

The minimum pressure of the helium occurs at a volume of approximately $1.44\times10^{-2}\text{m}^3$, where the pressure is approximately $2.88\times10^5\text{Pa}$ and the temperature is exactly 500K.

Sixteen miscellaneous worked examples lie behind you. The Digging Deeper questions that follow are now yours to attempt.

DIGGING DEEPER EXERCISE 5

EXERCISE 5A: BINDER QUESTIONS

Question 1

Two students write the first law of thermodynamics in different forms. One writes $\Delta U = Q - W$; the other writes $\Delta U = Q + W$. Each insists that the other is wrong. Resolve the disagreement by identifying what W means in each convention, and account for why the two equations are equivalent statements of the same physics.

Question 2

During an isothermal compression of an ideal gas, heat must leave the gas if its temperature is to remain constant. A student concludes from this fact that the internal energy of the gas must therefore decrease. Identify the error in the student's reasoning, and account for what becomes of the work that the surroundings do on the gas during the compression.

Question 3

Two different processes carry an ideal gas from the same initial state A to the same final state B along different paths on a PV diagram. Account for the fact that the change in internal energy ΔU is identical for both processes but the work done W and the heat exchanged Q can differ between them.

Question 4

Heat-capacity tables quote a single value of specific heat capacity for solids and liquids, but two distinct molar heat capacities (at constant volume and at constant pressure) for gases. A student concludes that the inequality $C_p > C_v$ is a feature peculiar to gases and does not exist for solids or liquids. Justify the student's observation by accounting for what changes, and what does not change, when energy is added to each phase of matter at constant pressure.

Question 5

Two ideal gases A and B are observed in separate experiments to have the same molar heat capacity at constant volume C_v. A student concludes from this single piece of evidence that the two gases must be the same substance, or at least be molecules of the same type. Evaluate this conclusion: justify what equal C_v values do, and do not, tell you about A and B at the molecular level, and identify what additional measurement would let you distinguish A from B.

Question 6

A welded steel pressure vessel sits in a Dar es Salaam yard. Overnight the temperature drops from 32°C to 22°C. State, with reason from the equation of state of an ideal gas, whether the pressure inside the vessel falls by exactly $10/305$ of its initial value, by more than $10/305$, or by less than $10/305$, and identify the assumption that underlies your conclusion.

Question 7

Two identical balloons are filled, one with helium and the other with ordinary air, both at the same temperature and at atmospheric pressure. The helium balloon is found to be much lighter than the air balloon. Account, by reference to Avogadro's law, for why the lighter helium balloon nevertheless contains the same number of molecules as the heavier air balloon, and identify the molecular property responsible for the mass difference.

Question 8

A student claims that, since the temperature of an ideal gas does not change during an isothermal process, no heat can flow into or out of the gas. Identify the error in this reasoning, and account for what the heat flow accomplishes if it is not raising the temperature.

Question 9

A student claims that any process happening rapidly must be adiabatic in the sense that the relation $PV^{\gamma} = \text{constant}$ holds throughout. Evaluate this claim, distinguishing carefully between the condition $Q \approx 0$ and the further condition that the process is quasi-static, and identify which condition must be added to 'happens rapidly' for the relation $PV^{\gamma} = \text{constant}$ to apply.

Question 10

On a single PV diagram, sketch an isothermal compression and an adiabatic compression of the same ideal gas, both starting from the same initial state and reducing the volume by the same factor. Mark the two end-states clearly, indicate which of the two compressions reaches the higher final pressure, and account for the molecular reason behind the difference.

Question 11

A closed cycle on a PV diagram is traversed in the counter-clockwise sense. From this single piece of information alone, identify with reason: (a) the sign of the net work done by the gas over the complete cycle, (b) the sign of the net heat absorbed by the gas over the complete cycle, (c) whether the device is functioning as a heat engine or as a refrigerator.

Question 12

A heat pump installed to warm a Dar es Salaam house during cool July weather is observed to deliver, for each joule of electrical energy it consumes from the grid, more than one joule of heat to the interior of the house. Account, by reference to the first law of thermodynamics applied to the heat pump's working substance, for why this is not a violation of energy conservation.

Question 13

A student argues that the Kelvin-Planck and Clausius statements of the second law of thermodynamics cannot both be correct, because the first speaks about heat engines while the second speaks about the spontaneous flow of heat, and the two sound completely unrelated. Resolve the student's confusion by identifying the single underlying physical principle that both statements express, and account for why a violation of either statement would necessarily imply a violation of the other.

EXERCISE 5B: REAL QUESTIONS

Question 14

A factory technician at an aluminium plant in Dar es Salaam observes that the compressed air emerging from the plant's industrial compressors feels noticeably hot to the touch. The technician reasons that *'energy is conserved, so the heat absorbed by the gas during the compression must equal the work done on it by the compressor.'* Identify the flaw in the technician's reasoning, and describe an everyday observation in industrial machinery that demonstrates the flaw directly.

Question 15

A boiler-plant operator in Mwanza, briefing a junior technician at the end of a shift, describes *'the gas doing 5000J of work'*, without specifying whether the work is done by the gas on the surroundings or by the surroundings on the gas. Account for why this ambiguity matters for first-law bookkeeping, and describe how a careful technician would record the same physics unambiguously.

Question 16

During a road accident, a sealed LPG cooking-gas cylinder is thrown into the wreckage of a burning vehicle. Account, by reference to the first law of thermodynamics, for why such cylinders are at risk of catastrophic explosion despite being designed for high-pressure operation. State the safety feature that manufacturers add to convert this risk from explosion to controlled release, and describe in one sentence how the feature works.

Question 17

A mother at home heats milk in an open saucepan and notices that the milk surface rises slightly as it warms, well before any boiling begins. Account for this rise in terms of an isobaric thermal expansion, and explain why the rise stops being noticeable once the milk reaches its boiling point.

Question 18

A bicycle pump becomes warm to the touch after vigorous use, even though no flame or external heater has been near it. A cyclist suggests that friction in the pump's seal is responsible for the warming. A physicist friend insists that this is at most a secondary effect and that the dominant cause is something else. Reconstruct the physicist's argument, and justify by appeal to direct observation why the friction explanation alone cannot account for the warming.

Question 19

Kipanga reads in a textbook that the indicator diagram of an internal combustion engine can be approximated by two adiabats and two isobars on a PV diagram, traversed clockwise. He asks **Mr. Akilikubwa** why all four cornerstone thermodynamic processes (isochoric, isobaric, isothermal, and adiabatic) are not used equally in the cycle, given that the chapter has spent so much effort on each. Provide the explanation Mr. Akilikubwa would give, framed in terms of the mechanical and thermal constraints of a real engine running at thousands of cycles per minute.

Question 20

A technician in Arusha claims to have built a 'self-powered fan' that runs on warm air alone, with no electrical input or fuel, by extracting energy from the heat of the surrounding atmosphere. Identify which form of the second law of thermodynamics this proposed device would violate, and account for why the device is impossible no matter how cleverly it is engineered.

Question 21

A government official in the energy ministry proposes building a coal-fired power station with a thermal efficiency of 80%, operating between a boiler at 500K and a cooling tower at 300K. Evaluate the proposal by reference to Carnot's theorem, and state with reason whether the project is feasible as proposed. If it is not, briefly suggest one direction in which the design parameters would have to move to make 80% theoretically achievable.

Question 22

An LPG cooking-gas cylinder used in a Dodoma household is moved from a shaded outdoor storage area at 25°C into a hot indoor kitchen at 40°C, with the valve closed throughout. The cylinder's safety pressure-release valve was calibrated to open when the internal pressure exceeds the rated limit, on the assumption that the cylinder is at 25°C. Account for the physical change inside the cylinder as it is brought indoors, and discuss whether the safety calibration is now conservative (over-cautious), exact, or risky.

Question 23

A factory in Tanga installs a refrigeration system to cool its server room. The factory engineer informs the manager that *'the refrigerator dumps more heat into the surroundings than it absorbs from the server room itself.'* A new accountant, suspicious of this claim, alleges that it violates the conservation of energy. Resolve the accountant's confusion by applying the first law of thermodynamics to the refrigeration cycle.

Question 24

A laboratory technician in Babati compresses 2.0mol of an ideal gas isothermally at 300K from a volume of 4.0L to a volume of 1.0L, and observes that approximately 6.9kJ of heat must be removed from the gas during the compression to keep its temperature constant. (a) State, with reason, whether the change in internal energy of the gas during the compression is positive, negative, or zero. (b) Account, by reference to the first law of thermodynamics, for why heat must be removed from the gas even though its temperature does not change. (c) Verify the magnitude of the heat removed by computing the work done on the gas. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 25

A pressure cooker boils food more rapidly than an open pan of water because food cooks faster at higher temperatures. Explain, by reference to the relationship between the saturation vapour pressure of water and its boiling temperature, why a sealed cooker enables higher cooking temperatures than an open pan, and account for the role of the cooker's safety valve in preventing the internal pressure (and hence the cooking temperature) from rising indefinitely.

EXERCISE 5C: HOT QUESTIONS

Useful constants: $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$**, atmospheric pressure** $= 1.013 \times 10^5\text{Pa}$**.**

Question 26

An ideal monatomic gas is taken from state A at (P_0, V_0, T_0) to state B by two different paths. Path 1: isothermal expansion from A until the volume is $2V_0$, then isochoric heating until the pressure returns to P_0. Path 2: isobaric expansion from A at pressure P_0 until the volume is $2V_0$. (a) Show that both paths terminate at the same final state B. (b) Verify that ΔU is the same along both paths. (c) Find the difference in heat absorbed between path 2 and path 1, expressed in terms of P_0V_0.

Question 27

An ideal gas of molar mass 16gmol^{-1} is found to have $\gamma = 1.30$ at room temperature. Identify the most likely number of accessible degrees of freedom for the gas, name a real gas consistent with these properties, and justify the identification.

Question 28

An ideal gas at pressure P_0, volume V_0, and temperature T_0 undergoes a quasi-static adiabatic compression to volume $V_0/4$. The gas has $\gamma = 1.4$. Find the final pressure and the final temperature, both as multiples of the initial values, and find the work done on the gas as a multiple of P_0V_0.

Question 29

The compression stroke of a four-stroke petrol engine takes about 1ms. Above a 5.0mm-thick aluminium piston of cross-section 0.005m², 0.5g of gas with specific heat capacity $1000\text{Jkg}^{-1}\text{K}^{-1}$ sits at a temperature 100K above the piston. Taking the thermal conductivity of aluminium as $230\text{Wm}^{-1}\text{K}^{-1}$, estimate the time for heat to leak from the gas through the piston by conduction, and by comparing the two timescales justify whether the adiabatic approximation is appropriate for the compression stroke.

Question 30

A heat engine using an ideal monatomic gas operates on a rectangular cycle on a PV diagram with corners at $A(V_1, P_1)$, $B(V_1, P_2)$, $C(V_2, P_2)$, $D(V_2, P_1)$, traversed in the order $A \to B \to C \to D \to A$, with $V_1 = 1.0 \times 10^{-3}\text{m}^3$, $V_2 = 4.0 \times 10^{-3}\text{m}^3$, $P_1 = 1.0 \times 10^5\text{Pa}$ and $P_2 = 4.0 \times 10^5\text{Pa}$. Find the net work done per cycle, the heat absorbed per cycle, and the efficiency of the engine. Compare the efficiency with the Carnot ceiling between the two temperature extremes of the cycle.

Question 31

A diesel engine compresses air from atmospheric pressure and 27°C to one-sixteenth of its initial volume in a single quasi-static adiabatic compression stroke. The autoignition temperature of diesel fuel in compressed air is approximately 480°C.

Verify that the compression alone raises the air's temperature above the autoignition temperature, and explain why this makes the diesel engine self-igniting without the spark plug needed in a petrol engine. Take $\gamma = 1.4$ for air.

Question 32

An ideal gas at 76cmHg and 27°C is contained in a cylinder fitted with a heavy frictionless piston. The piston is suddenly released and the gas expands adiabatically until its volume doubles. Taking $\gamma = 1.4$, find the final pressure of the gas, and the final temperature.

Question 33

A Carnot heat engine operates between two reservoirs. With the cold reservoir at 300K and the hot reservoir at 500K, the engine produces 200J of work per cycle. The cold reservoir is now cooled to 250K, with the hot reservoir unchanged, and the engine continues to absorb the same quantity of heat from the hot reservoir per cycle. Find the new work output per cycle, and the percentage increase in efficiency.

Question 34

An ideal gas with $\gamma = 1.4$, initially at $1.0 \times 10^5 \text{Pa}$, 0.020m^3, and 300K, undergoes a two-stage process. Stage 1: quasi-static adiabatic compression to one-quarter of its initial volume. Stage 2: isobaric cooling at the resulting pressure until the temperature returns to 300K. Find the total work done on the gas across both stages, the heat rejected during the cooling stage, and the net change in internal energy at the end.

Question 35

A vertical glass tube of uniform cross-section is closed at the bottom and open at the top. A column of mercury of length 0.20m is trapped above a column of air of length 0.40m at 27°C. Atmospheric pressure is 76cmHg. The tube is now slowly inverted so that it is closed at the top and open at the bottom, with the temperature held constant at 27°C. Find the new length of the trapped air column.

EXAMINATION QUESTIONS FOR PART ONE

Useful constants: *Young's modulus of steel* $= 2.0 \times 10^{11} Pa$, *Young's modulus of copper* $= 1.1 \times 10^{11} Pa$, *Young's modulus of aluminium* $= 7.0 \times 10^{10} Pa$, $c_{water} = 4200 Jkg^{-1}K^{-1}$, $L_f(ice) = 3.36 \times 10^5 Jkg^{-1}$, $L_v(water) = 2.26 \times 10^6 Jkg^{-1}$, $R = 8.314 Jmol^{-1}K^{-1}$, *atmospheric pressure* $= 1.013 \times 10^5 Pa$, $\sigma = 5.67 \times 10^{-8} Wm^{-2}K^{-4}$, *Wien displacement constant* $b = 2.898 \times 10^{-3} mK$, $0°C = 273K$, *solar constant* $S = 1370 Wm^{-2}$ $k_B = 1.381 \times 10^{-23} JK^{-1}$, $N_A = 6.022 \times 10^{23} mol^{-1}$

Question 1

(a) A student claims that the strain energy stored in a stretched wire is the kinetic energy of its atoms. Identify the error and account for the form in which the energy is actually stored.

(b) A daladala's rear leaf-spring suspension creaks under heavy passenger load but is silent when the vehicle is empty. Account for the difference, with reference to the elastic deformation of the leaf springs.

(c) A steel wire of length 2.0m and cross-sectional area $1.5 \times 10^{-6}\text{m}^2$ is stretched by an axial force of 300N. Take Young's modulus of steel as 2.0×10^{11}Pa. Calculate the extension of the wire and the elastic strain energy stored in it.

Question 2

(a) A student examining a stress-strain curve for a metal identifies the proportional limit and the elastic limit as the same point. Identify the error and account for the physical difference between the two limits.

(b) A bicycle pump used in a **Kigoma** shop, after months of daily use, develops a permanent slight bend in its piston rod even though no single push has snapped the rod. Account for the bend.

(c) A copper wire of yield stress 2.5×10^8Pa and breaking stress 4.0×10^8Pa has a circular cross-section of diameter 1.0mm. Calculate the maximum load the wire can support before plastic deformation begins, and the maximum load it can support before fracture.

Question 3

(a) Account for why the strain energy per unit volume of an elastic material is given by $\frac{1}{2} \times \text{stress} \times \text{strain}$, and not by $\text{stress} \times \text{strain}$ alone.

(b) A child in **Shinyanga** builds a slingshot using a thick rubber band. He notices that stretching the band to twice its initial extension launches the projectile to approximately four times the previous range, at the same angle. Account for this observation.

(c) A spring of force constant 200Nm^{-1} is stretched 0.30m from its natural length and then released, propelling a 0.20kg projectile horizontally from rest along a frictionless surface. Calculate the elastic potential energy stored in the spring and the launch velocity of the projectile.

Question 4

(a) For a fluid, distinguish between the isothermal bulk modulus and the adiabatic bulk modulus, and account for why the two values differ.

(b) The speed of sound in fresh water at room temperature is observed to be substantially higher than the speed of sound in air at the same temperature. Identify the physical property of water mainly responsible, and account for why this property is so much larger in water than in air.

(c) An ideal monatomic gas is at a pressure of 1.5×10^5Pa. Calculate (i) the isothermal bulk modulus of the gas, (ii) the adiabatic bulk modulus of the gas. Take $\gamma = 5/3$.

Question 5

(a) A bar of varying cross-section carries the same axial tensile force throughout its length. Account for why the strain in the bar is not the same everywhere, and identify where in the bar the strain is greatest.

(b) A glass tumbler dropped on a tiled floor in a kitchen typically shatters into many fragments, while a metal cup dropped from the same height suffers only a small dent. Account for this difference in behaviour.

(c) A steel rod of uniform cross-sectional area A_0 carries an axial tensile force F. A small circular hole drilled through the rod reduces its local cross-sectional area to $A_0/4$ at the hole. Treating the stress concentration factor at the edge of the hole as 3, calculate the factor by which the peak local stress at the hole exceeds the stress in the unaffected part of the rod.

Question 6

(a) A cable made of n parallel wires of identical cross-section is loaded with a force F. A student claims that each wire carries the same load F because they share the load. Identify the error in the student's reasoning and state the correct sharing of the load among the n wires.

(b) The lift cables of a passenger lift in a **Dar es Salaam** office building consist of multiple parallel steel wire ropes rather than a single thicker rope. Account for two engineering reasons (beyond the simple multiplication of cross-section) why a multi-rope arrangement is preferred.

(c) A rigid horizontal beam is supported in equilibrium by three vertical wires of identical material and identical length, attached to the beam at equal spacing. The two outer wires each have cross-sectional area A; the middle wire has cross-sectional area 2A. A vertical load W is hung from the centre of the beam. Calculate the tension in each of the three wires.

Question 7

(a) A small mass attached to the lower end of a vertical elastic wire and pulled downward from its equilibrium position oscillates after release. Identify the type of motion that results, and account for why the period of this motion depends on the elastic properties of the wire as well as the mass.

(b) A bus driver in **Mtwara** notices a steady low-pitched hum from his bus's wheel-and-axle assembly at one particular cruising speed, but not at higher or lower speeds. Account for this resonance in terms of the elastic properties of the suspension.

(c) A 2.0kg mass hangs in equilibrium from a vertical steel wire of length 1.5m and cross-sectional area $5.0 \times 10^{-7}\text{m}^2$. Take Young's modulus of steel as $2.0 \times 10^{11}\text{Pa}$. The mass is pulled a small distance below its equilibrium position and released. Calculate the period of the resulting vertical oscillation.

Question 8

(a) A student plots stress against strain for a metal wire and finds that the curve becomes non-linear shortly after leaving the origin. Identify what physical limit has been passed, and account for why the slope of the straight portion alone gives the Young's modulus.

(b) A factory in **Mwanza** tensile-tests every batch of incoming steel wire before installing it in a suspension bridge. Account for why a tensile test alone is sufficient to detect a sample of below-specification steel, but is not sufficient to confirm an above-specification sample as fully usable.

(c) In a tensile test on a steel rod of original length 200mm and original diameter 10mm, the following data are recorded at the proportional limit: applied force 14kN; extension 0.180mm. Calculate the Young's modulus of the steel.

Question 9

(a) A student claims that a load suddenly released onto an elastic wire produces the same maximum extension as the same load placed onto the wire gradually. Identify the error and account for what happens to the wire in the moments after a sudden release.

(b) A worker on a **Sumbawanga** construction site lowers a sack of cement onto a wire-supported hook by carefully releasing the sack at the level of the hook, rather than dropping it from above his shoulders. Account for why the dropping method is more likely to snap the supporting wire even though the cement weight is the same in both cases.

(c) A 5.0kg load is released from rest at a height of 0.10m above the lower end of a vertical steel wire of original length 2.0m and cross-sectional area $1.0 \times 10^{-6}\text{m}^2$. Take Young's modulus of steel as $2.0 \times 10^{11}\text{Pa}$ and $g = 9.81\text{ms}^{-2}$. Calculate the maximum dynamic extension of the wire and the factor by which it exceeds the static extension produced by the same load.

Question 10

(a) The fundamental frequency of a stretched wire fixed at both ends depends on its tension, length, and linear mass density. Account for why a wire of higher Young's modulus does not, by that fact alone, have a higher fundamental frequency than a similar wire of lower Young's modulus.

(b) A guitarist tuning her instrument in **Zanzibar** finds that tightening the tuning peg of a string raises its pitch. Account for this behaviour by reference to the formula for the fundamental frequency of a stretched wire.

(c) A steel wire of length 0.75m and linear mass density $8.0 \times 10^{-3}\text{kgm}^{-1}$ is stretched between two fixed pegs and is observed to vibrate at its fundamental frequency of 220Hz. Calculate the tension in the wire.

Question 11

(a) A student claims that any physical property of a substance that varies with temperature can serve as the basis of a thermometer. Identify three additional requirements that the property must satisfy, and account for why each is needed.

(b) A laboratory in **Mpanda** maintains a working thermometer for routine use and a reference thermometer kept locked in a cupboard, used only to recalibrate the working thermometer. Account for why the reference thermometer is kept specifically for this purpose rather than simply rotating two identical thermometers.

(c) A platinum resistance thermometer has resistance 25.0Ω at the ice point and 35.5Ω at the steam point. In a constant-temperature bath, the resistance is 31.5Ω. Calculate the temperature of the bath on the platinum-resistance scale.

Question 12

(a) Account for why the constant-volume gas thermometer is regarded as the near-ideal thermometer over a wide range of temperatures.

(b) An industrial calibration laboratory in **Dar es Salaam** keeps a constant-volume gas thermometer as its primary reference standard, using it once or twice a year to recalibrate platinum resistance thermometers used for daily measurements. Account for why the gas thermometer, despite its bulk and inconvenience, is preferred as the primary reference.

(c) A constant-volume gas thermometer has a pressure of 0.800×10^5Pa when its bulb is at the triple point of water (273.16K). When the bulb is placed in a heated bath, the pressure becomes 1.099×10^5Pa. Calculate the temperature of the bath on the ideal-gas scale, in kelvin and in degrees Celsius.

Question 13

(a) A student claims that since a platinum resistance thermometer gives the same reading as the ideal-gas scale at the ice point and at the steam point, the two scales must agree at every temperature in between. Identify the error in this claim, and account for the physical origin of any disagreement.

(b) A power-station engineer uses a platinum resistance thermometer (rather than a mercury-in-glass thermometer) to monitor the temperature of superheated steam at about 400°C. Account for two reasons why the platinum thermometer is preferred for this application.

(c) A platinum resistance thermometer has resistance 25.0Ω at the ice point and 35.6Ω at the steam point. In a heated bath, its resistance is 41.6Ω. A constant-volume gas thermometer placed in the same bath gives a temperature reading of 158.0°C. Calculate the temperature of the bath on the platinum-resistance scale, and the discrepancy between this value and the ideal-gas scale value.

Question 14

(a) A thermocouple consists of two different metal wires joined at two junctions, one at a known reference temperature and the other at the unknown temperature to be measured. Account for the origin of the thermoelectric EMF generated between the two junctions, and identify why the cold junction is necessary.

(b) A meat-processing factory in **Iringa** uses a thermocouple probe (rather than a glass thermometer) to measure the temperature inside cuts of meat as they leave the oven. Account for two practical advantages of the thermocouple in this application.

(c) A copper-constantan thermocouple has its cold junction held in a bath of melting ice at 0°C. When the hot junction is held at the steam point of water (100°C), the EMF measured across the open ends of the thermocouple is 4.10mV. Treating the EMF as varying linearly with the temperature of the hot junction in the working range 0 to 200°C, calculate the temperature of an oil bath that produces an EMF of 7.40mV in the same thermocouple.

Question 15

(a) Account for why the modern thermodynamic temperature scale is defined using the triple point of water as a single fixed reference point, rather than the ice point and the steam point used historically.

(b) A standards laboratory in **Dar es Salaam** realises a triple-point cell of water as part of its temperature-reference chain. Identify what is special about a triple-point cell that makes its temperature more reproducible than a simple ice-water bath.

(c) A constant-volume gas thermometer reads a pressure of 100kPa when its bulb is at the triple point of water (273.16K). The same thermometer, with its bulb placed in a vessel of liquid nitrogen at its boiling point, reads a pressure of 28.0kPa. Calculate the temperature of the boiling point of liquid nitrogen on the ideal-gas scale, in kelvin and in degrees Celsius.

Question 16

(a) Account for why a thermistor (a semiconductor temperature sensor) has an electrical resistance that decreases with rising temperature, the opposite of a metal resistance thermometer, and account for why this property makes thermistors particularly suited to detecting small changes in temperature near room temperature.

(b) A digital clinical thermometer used in a **Bagamoyo** health clinic has a thermistor at the tip of its probe rather than a platinum resistance element. Account for two practical reasons (beyond cost) why a thermistor is preferred in this application.

(c) A thermistor has resistance 10.0kΩ at 25°C and 3.30kΩ at 50°C. The resistance varies with absolute temperature according to $R = R_0 e^{(B/T)}$, where R_0 and B are constants. Calculate the value of B and the resistance of the thermistor at body temperature 37°C.

Question 17

(a) Account, by reference to the physical structure of a clinical mercury-in-glass thermometer, for why the reading remains at the maximum reached during measurement, even when the thermometer has been removed from the patient and cooled to room temperature.

(b) A nurse at a **Chalinze** health clinic must shake the clinical thermometer vigorously after each patient's measurement before the next reading. Account for the physical reason this shaking is necessary.

(c) The bulb of a clinical thermometer contains 0.20cm³ of mercury at 0°C. The cubic expansion coefficient of mercury is $\gamma_{Hg} = 1.82\times10^{-4}\text{K}^{-1}$. Take the cubic expansion of the glass as negligible. Calculate the change in the volume of mercury when the thermometer reading rises from room temperature 24°C to the patient's body temperature 39.5°C.

Question 18

(a) Distinguish between the operating ranges and sensitivities of a mercury-in-glass thermometer and an alcohol-in-glass thermometer, and account for why each is preferred for its own range of applications.

(b) A weather station outside **Mbeya** in the Southern Highlands, where night-time temperatures occasionally fall below -20°C in the cooler months, uses an alcohol-in-glass thermometer (rather than a mercury-in-glass thermometer) for its outdoor minimum-temperature record. Account for the choice.

(c) A mercury-in-glass thermometer has a bulb containing 0.50cm³ of mercury at 0°C. It is graduated such that a 1.0mm rise of the mercury thread in its stem corresponds to a temperature rise of 1.0K. Take the cubic expansion coefficient of mercury as $\gamma_{\mathrm{Hg}} = 1.82 \times 10^{-4}\mathrm{K}^{-1}$ and assume the expansion of the glass is negligible. Calculate the cross-sectional area of the bore of the stem.

Question 19

(a) A radiation thermometer (pyrometer) measures the temperature of a body without making physical contact with it. Account for the physical principle that allows this measurement, and identify the upper temperature limit above which contact thermometers (mercury-in-glass, platinum resistance) become impractical or impossible.

(b) A steel-mill furnace operator at **Tanga** monitors the temperature of molten steel from outside the furnace using an optical pyrometer, while a poultry-incubator operator at a **Bagamoyo** farm uses a digital probe thermometer to monitor the temperature inside the incubator. Account for why each operator's choice of thermometer is appropriate to the application.

(c) An optical pyrometer is calibrated by viewing a black body at a series of known temperatures, and is found to give an output reading proportional to the fourth power of the absolute temperature, in agreement with the Stefan-Boltzmann law. When pointed at a black body at 1000K, the pyrometer reads 1.00 (in arbitrary units). When pointed at a different glowing body, treated as a black body, the pyrometer reads 4.00. Calculate the temperature of this body.

Question 20

(a) Distinguish between thermal conductivity and thermal conductance, and account for why the conductance of a wall containing a window of poorer thermal insulation than the wall itself depends not only on the conductivity of the two materials but also on the area of each.

(b) An office building in **Lindi** has a south-facing wall containing both a brick portion and a single glass window. The architect adds a second window of the same area on the same wall. Account for whether the heat lost through the wall doubles, less than doubles, or more than doubles, identifying the dominant heat-loss path.

(c) A wall of total area 10m² has a brick portion of area 8.5m² and thickness 0.10m, with thermal conductivity $1.0\mathrm{Wm^{-1}K^{-1}}$. A glass window of area 1.5m² and thickness 5.0mm fills the rest, with thermal conductivity $0.80\mathrm{Wm^{-1}K^{-1}}$. When the inside is at 25°C and the outside at 15°C, calculate the rate of heat conduction through (i) the brick portion, (ii) the glass window, and (iii) the wall as a whole.

Question 21

(a) Account for why a cylindrical hot-water pipe loses heat at a rate per unit length that depends, for a given temperature difference between the pipe surface and the surroundings, on the natural logarithm of the ratio of the outer to the inner radius of any insulating layer wrapped around it, rather than on the layer thickness alone.

(b) A solar water heater installed on a **Mwanza** rooftop carries hot water from the rooftop tank to a tap on the ground floor through a copper pipe wrapped in two distinct layers of lagging: an inner layer of mineral wool (low thermal conductivity) and an outer layer of foil-faced foam (slightly higher conductivity but also low). Account for why the engineer uses two layers in series rather than the same total thickness of a single material.

(c) A copper hot-water pipe of inner radius 2.0cm carries water at 80°C through a building where the surrounding air is at 20°C. The pipe is lagged with two concentric layers: an inner layer of thickness 2.0cm and thermal conductivity $0.05\mathrm{Wm^{-1}K^{-1}}$, surrounded by an outer layer of thickness 1.0cm and thermal conductivity $0.10\mathrm{Wm^{-1}K^{-1}}$. Treating the pipe surface as held at 80°C and the outermost lagging surface as held at 20°C, calculate the rate of heat loss per unit length of the pipe.

Question 22

(a) Newton's law of cooling, in the form rate of cooling proportional to temperature difference between the body and its surroundings, holds rigorously for forced convection but only as an approximation for natural convection at small temperature differences. Account for why this distinction arises.

(b) A barbecue cook in **Bukoba** finds that a piece of grilled meat cools faster if she fans it with a banana leaf than if she leaves it on the table to cool in still air. Account for the difference, identifying which heat-transfer mechanism is enhanced by the fanning.

(c) A small body cools from 80°C to 60°C in still air at 20°C in 5.0 minutes. When the same body is placed in a steady draught at the same air temperature, it cools from 80°C to 60°C in 1.5 minutes. Calculate the ratio of the cooling-law constant under forced convection to the constant under free convection.

Question 23

(a) A blackbody is heated continuously from 1000K through 2000K to 5000K. Account for the systematic shift in the wavelength at which it radiates maximum energy, and account for the corresponding change in the colour at which the body appears to a human observer.

(b) Astronomers in a Tanzanian observatory classify stars by colour: red giants appear cool red, our Sun appears yellow-white, and certain hot young stars appear bluish white. Account for this colour-temperature relationship by reference to Wien's displacement law.

(c) Wien's displacement constant is $b = 2.898 \times 10^{-3}\text{mK}$. A particular star emits its maximum radiation at a wavelength of 480nm. Calculate the surface temperature of the star, and compare it with the surface temperature of the Sun (which emits its maximum radiation at approximately 500nm).

Question 24

(a) A grey body has the same emissivity at all wavelengths, but the emissivity is less than that of a perfect blackbody at the same temperature. Account, by reference to Kirchhoff's law of thermal radiation, for why a body that is a poor emitter at a particular wavelength must also be a poor absorber at the same wavelength.

(b) The roof of a corrugated-iron house in a **Dar es Salaam** neighbourhood is painted matt black to reduce the temperature of the interior on a hot afternoon. Account, by reference to the difference in dominant radiation wavelengths during the day and at night, for whether the matt black paint actually achieves the desired cooling.

(c) A grey-bodied object of emissivity 0.60 is held at a steady temperature of 400K in surroundings at 300K. Take the Stefan-Boltzmann constant as $\sigma = 5.67 \times 10^{-8}\text{Wm}^{-2}\text{K}^{-4}$. Calculate the net rate of radiation loss per unit area of the object's surface.

Question 25

(a) A vacuum flask is designed to minimise heat transfer between its contents and the surroundings by all three mechanisms (conduction, convection, and radiation). Identify the design feature that minimises each mechanism, and account for why the flask cannot completely eliminate any of them.

(b) A market trader in **Chakechake** fills a vacuum flask with hot tea at 90°C in the morning and finds that the tea is still pleasantly hot at 70°C when she drinks it eight hours later. Account for the tea's slow cooling, and identify the residual heat-loss mechanism that ultimately limits the flask's effectiveness.

(c) A vacuum flask of internal surface area 0.010m² holds tea at 80°C in a room at 20°C. The internal walls of the flask are silvered, with effective emissivity 0.05. Take the Stefan-Boltzmann constant as $\sigma = 5.67 \times 10^{-8}\text{Wm}^{-2}\text{K}^{-4}$. Calculate the rate of heat loss by radiation through the silvered vacuum gap, and compare it with the approximate rate of heat loss from an unlagged open cup of the same area by free convection alone, taking the convection coefficient as $h = 10\text{Wm}^{-2}\text{K}^{-1}$.

Question 26

(a) A solar absorber panel is painted matt black on its sun-facing side to absorb the maximum fraction of incident solar radiation, but the same blackness causes it to radiate strongly in the infrared at its operating temperature. Account for why this radiative loss is unavoidable, and identify a coating strategy (a 'selective surface') that allows a high absorption of solar radiation without the corresponding high emission of thermal radiation.

(b) A solar water heater installed on a rooftop has a flat absorber panel of area 2.0m². On a clear day at noon the panel receives a solar flux of 800Wm⁻² and reaches an operating temperature of 70°C while the surroundings are at 30°C. Account for why the panel does not continue to heat indefinitely, and identify the limiting factor on its steady-state temperature.

(c) A solar absorber panel of area 2.0m² and emissivity 0.90 (assumed equal to its solar absorptivity) sits on a rooftop and receives solar radiation of flux 800Wm⁻². At its steady operating temperature of 70°C in surroundings at 30°C, calculate the rate at which the panel absorbs solar radiation, the rate at which it loses heat by infrared radiation to the surroundings, and the net useful heat-collection rate.

Question 27

(a) Account, by reference to thermal radiation to the sky, for why frost can form on the surface of a tin roof or on the leaves of a plant on a clear night even when the air temperature one metre above the surface is several degrees above 0°C.

(b) A farmer in **Mufindi** observes that on clear cold nights his crops in an open field develop frost on their leaves, while on cloudy nights with the same air temperature no frost appears. Account for the role of cloud cover in suppressing frost formation.

(c) On a clear night, the effective sky temperature for thermal radiation is approximately 250K, while on an overcast night with low clouds it is approximately 280K. The ground surface is at 5°C (278K) on both nights and emits as a blackbody. Calculate the net rate of radiative heat loss per unit area of ground surface for each case, and account for why one case can produce frost while the other cannot.

Question 28

(a) The Earth, treated as a perfect absorber of solar radiation and a perfect emitter of thermal radiation, would settle at a steady-state temperature determined by balancing the absorbed solar power with the thermal-infrared emission to space. Account for the geometric factor (involving the Earth's cross-section as seen from the Sun and its surface area for emission) that enters this balance.

(b) Climate scientists at a Tanzanian research institute observe that the average measured temperature at the Earth's surface (around 288K) is substantially higher than the calculated equilibrium temperature for a bare planet of the same albedo. Account for this discrepancy by reference to the Earth's atmosphere.

(c) Take the solar constant at the Earth's orbit as $S = 1370\text{Wm}^{-2}$ and the Earth's average albedo (the fraction of solar radiation reflected back to space) as 0.30. Treating the Earth as a sphere that absorbs the un-reflected solar radiation and re-emits as a blackbody, calculate the equilibrium effective temperature of the Earth on the Stefan-Boltzmann scale.

Question 29

(a) A typical horticultural greenhouse uses glass walls and roof to admit solar radiation while restricting the loss of heat at night. Account for the optical property of glass that allows it to act in this dual role, and identify the wavelength region in which glass is transparent and the region in which it is opaque.

(b) A market gardener in **Iringa** builds a polythene-sheet greenhouse to extend the growing season for tomatoes. The polythene is transparent to visible light but, unlike glass, is also fairly transparent to thermal infrared. Account for why the polythene greenhouse is less effective at retaining heat than a glass one, and account for the limited improvement that polythene provides over no enclosure at all.

(c) A greenhouse glass panel of area 4.0m² is at a steady temperature of 20°C while the air outside is at 5°C. The glass behaves as a blackbody for thermal infrared radiation. Calculate the rate at which the panel loses heat to the cooler outdoor air by net infrared radiation, comparing this with the rate at which an equivalent perfectly transparent (zero-emissivity) film of the same area would lose. (Take the air's effective radiation temperature as 5°C.)

Question 30

(a) Account, by reference to the kinetic theory of an ideal gas, for why the root-mean-square speed of the molecules of a gas at a given temperature depends on the molar mass of the gas, but the average translational kinetic energy per molecule does not.

(b) At a balloon-display event, two helium-filled balloons are released simultaneously with hydrogen-filled balloons. The hydrogen balloons are observed to deflate (lose gas through the rubber wall) noticeably faster than the helium balloons, even though both gases are at the same temperature and the rubber walls are identical. Account for this observation, identifying the molecular property responsible.

(c) Calculate the root-mean-square speeds of hydrogen and nitrogen molecules at 300K, given molar masses $\text{M}_{\text{H}_2} = 2.0 \times 10^{-3}\text{kgmol}^{-1}$ and $\text{M}_{\text{N}_2} = 28 \times 10^{-3}\text{kgmol}^{-1}$. Find the ratio of the two speeds and comment briefly on its physical significance. Take $\text{R} = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 31

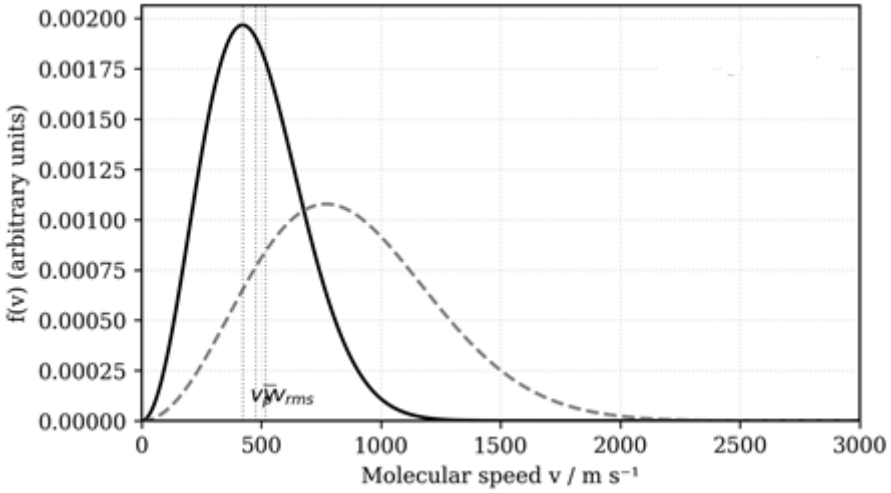
(a) Account for why the total translational kinetic energy of all the molecules in a sample of an ideal gas is far larger than the kinetic energy of a single sports ball, even though each individual molecule has only a tiny kinetic energy and even though the gas appears to be at rest.

(b) A football placed on a windless day at the centre of a **Benjamin Mkapa** pitch appears stationary. A nearby kicked football flies through the air. State, with reason from the kinetic theory of gases, whether the stationary football's mass-and-air system contains more, less, or comparable translational kinetic energy than the kicked football itself.

(c) Calculate the total translational kinetic energy of 2.0 mol of an ideal gas at 300K. Compare this with the kinetic energy of a 0.45kg football kicked at 25ms^{-1}. Take $\text{R} = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 32

The following figure shows the Maxwell-Boltzmann speed distribution for molecules of a fixed-mass sample of nitrogen at two different temperatures.



(a) Identify the curve corresponding to the higher temperature, and account for the shift of the peak and the broadening of the distribution as the temperature rises.

(b) On a hot afternoon in Dodoma, a few molecules of perfume sprayed in one corner of a still room are smelt at the opposite corner only some seconds later, even though the mean speed of perfume molecules at 300K is comparable to the speed of a fast bicycle. Account, by reference to the Maxwell-Boltzmann distribution and to the molecules' frequent collisions, for the apparent slowness of perfume diffusion.

(c) From the figure, identify the most probable speed and estimate the rms speed of nitrogen molecules at 300K. Compare your estimate of the rms speed with the calculated value from $v_{rms} = \sqrt{3RT/M}$, taking $M_{N_2} = 28 \times 10^{-3}\text{kgmol}^{-1}$ and $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 33

(a) Distinguish between effusion and diffusion as processes by which gas molecules escape from one region to another, and account for why Graham's law *(rate of effusion inversely proportional to the square root of molar mass*) holds for effusion but is only approximately applicable to diffusion.

(b) A factory in Tanzania uses a long set of porous-membrane stages to enrich uranium for nuclear-power use, where the natural mixture of UF_6 molecules contains both $^{235}UF_6$ and $^{238}UF_6$ at very similar molar masses (349 and 352 g/mol respectively). Account for why a single membrane stage gives only a tiny enrichment of the lighter isotope, and why many stages in series are required.

(c) Two identical containers each have a small hole in the wall. One contains helium gas, the other oxygen gas, both at the same pressure and temperature. Calculate the ratio of the rate at which helium effuses through its hole to the rate at which oxygen effuses through the identical hole. Take $M_{He} = 4.0 \times 10^{-3}\text{kgmol}^{-1}$, $M_{O_2} = 32 \times 10^{-3}\text{kgmol}^{-1}$.

Question 34

(a) The speed of sound in an ideal gas is given by $v = \sqrt{\gamma\, RT/M}$. Account for why a person who inhales helium and then speaks has an audibly higher-pitched voice than usual, even though the vocal cords themselves vibrate at the same frequency as before.

(b) A safety officer at a Tanga industrial gas plant warns workers against inhaling helium directly from a cylinder for amusement, despite the high-pitched-voice effect being apparently harmless. Account for the physical danger of inhaling pure helium, identifying the property of helium that the warning addresses.

(c) Calculate the speed of sound in dry air and in helium, both at 27°C, taking $\gamma_{air} = 1.40$, $M_{air} = 29 \times 10^{-3}\text{kgmol}^{-1}$, $\gamma_{He} = 5/3$, $M_{He} = 4.0 \times 10^{-3}\text{kgmol}^{-1}$, and $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$. Find the ratio of the two speeds.

Question 35

(a) A diatomic molecule at room temperature stores its thermal energy in five accessible degrees of freedom: three translational and two rotational. Account, by reference to the geometry of a dumbbell-shaped molecule, for why two and not three rotational modes are available.

(b) A factory in Mwanza compares the cost of filling its argon balloons and its nitrogen balloons. The technician notes that for the same volume and temperature, the gases contain the same number of moles, but the nitrogen balloon has more total internal energy than the argon balloon. Account for this observation in terms of degrees of freedom.

(c) Calculate the total internal energy of 1.0 mol of nitrogen (N_2) at 300K. Find the fraction of this internal energy that resides as translational kinetic energy and the fraction that resides as rotational kinetic energy. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 36

(a) The molar heat capacity at constant volume of an ideal gas with f accessible degrees of freedom is $C_v = (f/2)R$, and at constant pressure is $C_p = ((f+2)/2)R$. Account, by reference to the first law of thermodynamics applied at constant pressure, for the additional R in C_p relative to C_v.

(b) A welder in a workshop uses bottled argon gas as a shielding gas for arc welding because the arc stays cooler in argon than it would in nitrogen at the same heating power. Account for this observation in terms of the heat capacities of the two gases.

(c) For a diatomic ideal gas with $f = 5$ accessible degrees of freedom, calculate C_v, C_p, and the ratio of principal heat capacities $\gamma = C_p/C_v$. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 37

(a) The first law of thermodynamics states that the change in internal energy of a closed system equals the heat added to it minus the work done by it. Account for why, in any heating process at constant pressure, the heat supplied to an ideal gas always exceeds the change in its internal energy by exactly the work done by the gas during the expansion.

(b) A welder in a workshop fills a flexible balloon with helium from a cylinder, then leaves it on a sunny bench. The sun warms the helium and the balloon expands. Account, by reference to the first law of thermodynamics, for why the heat absorbed by the helium from the Sun exceeds the rise in the gas's internal energy by an amount equal to the work the helium has done in pushing back the surrounding atmosphere.

(c) 0.50 mol of an ideal monatomic gas at 300K is heated at the constant pressure $1.0 \times 10^5\text{Pa}$ until its temperature is 600K. Calculate the heat supplied, the work done by the gas, and the change in its internal energy. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 38

(a) The change in internal energy of an ideal gas between two given states depends only on the two states, not on the path joining them. Account for why the work done by the gas in moving between the same two states is different along different paths.

(b) A factory technician in Mwanza uses a piston compressor to raise a gas from atmospheric pressure to operating pressure. He notices that the compressor uses less work in cool weather (when the gas is allowed to remain near atmospheric temperature during compression) than in hot weather (when the gas heats during the compression and the pressure rises more rapidly). Account for this observation.

(c) An ideal gas undergoes a transition from state A at (P_0, V_0) to state B at $(P_0/2, 2V_0)$ along two different paths. Path 1: isochoric drop from A to $(P_0/2, V_0)$, then isobaric expansion to B. Path 2: isobaric expansion from A to $(P_0, 2V_0)$, then isochoric drop to B. Calculate the work done by the gas along each path, expressed in terms of P_0V_0, and identify which path delivers more work.

Question 39

(a) The net work done by an ideal gas during one complete traversal of any closed cycle on a PV diagram equals the area enclosed by the cycle, with the sign of the area set by the direction of traversal. Account, by reference to the integral $W = \oint P\,dV$, for why this geometric area rule applies to cycles of every shape, not just rectangular ones.

(b) A small two-stroke engine in a motorbike has its indicator diagram (a real-time PV record of the gas in the cylinder) recorded by a diagnostic technician. The diagram is roughly triangular rather than perfectly rectangular. Account for why the technician can still infer the engine's net work output per cycle from the shape of the diagram alone.

(c) An ideal gas executes a triangular cycle on a PV diagram with corners at $A(V = 1.0 \times 10^{-3}\text{m}^3, P = 1.0 \times 10^5\text{Pa})$, $B(V = 1.0 \times 10^{-3}\text{m}^3, P = 2.0 \times 10^5\text{Pa})$, and $C(V = 2.0 \times 10^{-3}\text{m}^3, P = 1.0 \times 10^5\text{Pa})$, traversed in the order $A \to B \to C \to A$. The path $B \to C$ is a straight line on the PV diagram. Calculate the net work done by the gas per cycle.

Question 40

(a) In an isothermal expansion of an ideal gas the temperature is held constant by allowing heat to flow in from the surroundings. Account, by reference to the first law and the property that internal energy of an ideal gas depends only on temperature, for why the heat absorbed during the expansion exactly equals the work done by the gas.

(b) A natural-gas pipeline in a Tanzanian gas-distribution network carries compressed methane from a wellhead at 80 atmospheres into a distribution network at 5 atmospheres, with the gas held at approximately ground temperature by long contact with the soil. Account for why the heat absorbed by the gas from the soil equals the work the gas does in expanding, even though the gas does not visibly heat up during the expansion.

(c) One mole of methane (treated as an ideal gas) at 290K is expanded isothermally from a pressure of 80 atmospheres to a pressure of 5 atmospheres. Calculate the work done by the gas during the expansion. Take $R = 8.314\text{Jmol}^{-1}\text{K}^{-1}$.

Question 41

(a) Carnot's theorem fixes the maximum efficiency of any heat engine operating between two reservoirs at temperatures T_{hot} and T_{cold}, in kelvin, as $\eta_{\text{Carnot}} = 1 - T_{\text{cold}}/T_{\text{hot}}$. Account for why this maximum cannot be approached except in the

limit of a reversible (and therefore infinitely slow) cycle, and account for the consequences for real engines run at finite power.

(b) A geothermal power station built near the volcanic Lake Natron in northern Tanzania extracts heat from underground steam at 150°C and rejects waste heat to the atmosphere at 27°C. Account, with reason from Carnot's theorem, for why the engineer designing the plant should expect a thermal efficiency of at most about 30%, and identify two factors that drive the realised efficiency below this Carnot ceiling.

(c) A heat engine operates between an underground steam reservoir at 150°C and a cooling tower at 27°C. Calculate the Carnot efficiency for this engine, and the maximum work that could be extracted from each kilojoule of heat absorbed from the steam reservoir.

Question 42

(a) In a calorimetry experiment, the measured final temperature is always slightly lower than the true value if the system is losing heat to the surroundings during the mixing process. Explain which mechanism is primarily responsible and why the error increases with longer experimental times.

(b) A domestic refrigerator manufacturer in Mwanza advertises its product as having a coefficient of performance of 12 between an interior temperature of 4°C and a kitchen temperature of 28°C. State, with reason from Carnot's theorem applied in reverse, whether this claim is feasible.

(c) A domestic refrigerator maintains its interior at 4°C in a kitchen at 28°C. Its compressor consumes 200W of electrical power. Calculate the maximum possible coefficient of performance for this refrigerator, and the maximum rate at which it could extract heat from its interior.

Question 43

(a) A corrugated iron roof in Dar es Salaam becomes extremely hot in midday sun but cools rapidly after sunset. Explain the rapid cooling, identifying which heat-transfer mechanisms dominate during the night-time cooling and why the roof cools faster than the concrete walls of the same building.

(b) A gas-filled cylinder in an industrial laboratory is briefly connected through a valve to an adjacent evacuated cylinder. Both cylinders are well insulated. The gas rushes from the filled to the empty cylinder until pressure equalises. Account for why the final temperature of the gas (an ideal gas, well above its boiling point) is essentially the same as its initial temperature, even though the expansion is rapid and the system is thermally insulated.

(c) 0.50 mol of an ideal gas at 27°C and 1.0×10^{5}Pa fills one half of a thermally insulated 2.0×10^{-3}m³ container, with the other half evacuated. The partition is suddenly removed. Calculate the final temperature, the final pressure, and the change in internal energy of the gas. Take $\mathrm{R} = 8.314\mathrm{J mol^{-1} K^{-1}}$.

PART TWO

VIBRATIONS AND WAVES

Having explored the behaviour of matter and the principles of heat in Part One, we now turn to motions that repeat and disturbances that travel through matter and space. Vibrations and Waves examines oscillatory motion such as simple harmonic motion, the propagation of mechanical waves, the nature of sound, the behavior of light in physical optics, electromagnetic waves, and the Doppler effect. These ideas explain how energy is transferred without the permanent movement of matter and help us understand systems ranging from vibrating strings and sound waves to optical instruments, radio transmission, and modern communication technologies. Mastering this part will deepen your understanding of periodic motion and wave behaviour, and prepare you for more advanced studies in modern physics and applied science.

Chapter 6

SIMPLE HARMONIC MOTION

INTRODUCTION

The trouble began, as it often does at Miono Secondary, on a Tuesday afternoon at the school sports field. Kipanga was attempting, with limited success, to push Kipute on the metal swing near the boundary wall. Each time she reached the lowest point and started moving forward, his push arrived a fraction too late, by which point she was already pulling away from him, and his hand bumped gently against the back of an empty seat. On the return swing he pushed too early. By the seventh attempt the swing was making the kind of half-hearted excursion a swing makes only when it suspects it is being mocked.

"Kipanga," Kipute said over her shoulder, in the tone she reserved for moments when his physics required correction, *"you are fighting the swing."*

"I am pushing the swing."

"You are pushing it at the wrong time. Wait until I reach the bottom and am moving away from you. Push then. Then let me come back. Then push at exactly the same moment of the next cycle. Match the swing's own rhythm. Stop arguing with it."

Kipanga waited. He pushed once, gently, at exactly the moment Kipute crossed the lowest point. He waited again. He pushed at the same moment of the next cycle. By the fifth push the swing was rising satisfyingly. By the tenth, Kipute was higher than the boundary wall and Kipanga was looking unreasonably proud of himself.

"Why," he asked the air, *"does it matter when I push?"*

Mr. Akilikubwa, who had a habit of materialising at boundary walls when good questions were imminent, materialised at the boundary wall. He had the contented look of a man who had been waiting twenty-three years for exactly this question.

"Because the swing has its own natural frequency," he said. *"It wants to oscillate at one particular rate, determined entirely by its length and the local strength of gravity. If you push at that rate, in resonance with its rhythm, every push adds a little energy and the amplitude grows. If you push at any other rate, you fight the swing and your energy is wasted, mostly into the muscles of your arm. Kipute discovered this by instinct. Newton discovered it by mathematics. The mathematics begins in this chapter."*

He walked over, caught the seat, and brought the swing gently to rest. *"This swing is the simplest oscillator in nature. One restoring force, one natural frequency, one rhythm. But the same mathematics governs a car suspension, a guitar string, a quartz crystal in your phone, the columns of air in your lungs when you speak, and the basilar membrane in your inner ear that lets you tell those columns apart. All of them are, in their hearts, the same thing. They obey one equation, and that equation rests on one observation: when you pull a thing away from where it wants to be, the harder you pull, the harder it pushes back. Force proportional to displacement, always pointing back home. That is* ***simple harmonic motion.*** *That is what this chapter is about."*

Welcome to Chapter 6. The previous five chapters spent themselves on matter that pushes back when you squeeze it (Chapter 1), matter that warms up when you heat it (Chapter 2), matter that conducts that warmth to its neighbours (Chapter 3), matter that keeps its temperature by the busy random motion of its molecules (Chapter 4), and matter that converts heat into work in heat engines and into nothing useful in everyday life (Chapter 5). All of that was about state and slow change. This chapter is the first to be about repetition. It is about systems that do the same thing, over and over, at the same rate, until they get tired and stop, and then start again the moment something gives them a nudge in time.

If you remember the molecular potential well from Chapter 1 (the picture in which two bonded molecules sit at the bottom of a curve that looks like a spring near its minimum), you already know why simple harmonic motion is universal. Anything that lives near the bottom of a smooth potential well, when displaced a little, behaves like a tiny mass on a tiny spring. The world is full of such wells, and this chapter is about the motion that happens inside them.

We also have an old promise to keep. Back in Chapter 1, when we discussed the shear modulus, an experimental method for measuring it was deferred: the one in which a wire is suspended from a rigid support, a horizontal disc is fixed at its lower end, and the disc is set into torsional oscillations. We could not write down the period of those oscillations at the time, because the period belongs to the language of simple harmonic motion. We will keep that promise in this chapter, in the section on torsional oscillations, where the answer comes out so cleanly it almost feels unfair.

And we will need a third tool from earlier in the series. The chapter on the rotation of rigid bodies introduced the moment of inertia, the parallel axis theorem, and the radius of gyration. Those ideas were treated as endpoints in their own right. Here, they become starting points. Three sections of this chapter, on the compound pendulum, the torsion pendulum, and the rolling oscillation in a bowl, could not be written without them. If those concepts are not fresh in your memory, take a moment to revisit the rotational dynamics chapter before you reach the second half of this one. The calculation style here is direct, but only if the moment of inertia is not a mystery.

THE OSCILLATION VOCABULARY

Before we write down a single equation, we need to agree on a small vocabulary. Almost every word in this chapter will refer to some feature of a body that is moving back and forth. If the words are slippery, the equations will feel slippery; if the words are sharp, the equations will feel inevitable. So we begin slowly.

Periodic motion and oscillation

A motion is **periodic** if *it repeats itself, exactly, after a fixed interval of time*. The motion of the hands of a clock is periodic. The orbit of the moon around the Earth is periodic. The bouncing of a ball on a hard floor, if no kinetic energy were lost, would be periodic. Periodic motion need **not** pass through the same position twice in opposite directions; the moon, for example, simply returns to where it began without ever reversing its direction of travel. Repetition is enough.

An **oscillation** is the special kind of periodic motion that involves *a back-and-forth excursion about a fixed central position*. A swing oscillates. A clock pendulum oscillates. A guitar string oscillates. A child on a trampoline, neglecting air resistance, oscillates. The moon does not oscillate. Every oscillation is periodic, but not every periodic motion is an oscillation. In this chapter the word *oscillation* will always mean a back-and-forth motion of this sort.

Equilibrium position, displacement, and amplitude

The fixed central position about which the body oscillates is called the **equilibrium position**. It is *the position the body would settle at if it were left alone with no nudge:* the bottom of the swing, the natural length of the spring, the still surface of a pond. In every oscillation we will study, the equilibrium position is the position at which the net force on the body is zero. Pull the body away from this position and a force appears pulling it back.

The **displacement**, written **x**, is *the signed distance of the body from its equilibrium position at any instant.* We choose one direction along the line of motion to be the positive direction; then x is positive on that side of equilibrium and negative on the other. Displacement is therefore a number that may be positive, negative, or zero, and which changes continuously as the oscillation proceeds. At the equilibrium position itself, $x = 0$.

The **amplitude**, written A, is *the **largest** value the displacement ever takes during an oscillation*. It is always a positive number. The body therefore moves between $x = +A$ and $x = -A$ on either side of equilibrium, and the **total length** of its excursion from one extreme to the other is therefore $2A$. For a pendulum we may use the angular amplitude (the largest angle from vertical) instead of a linear amplitude, depending on which is more convenient; the following figure shows both senses on the same diagram.

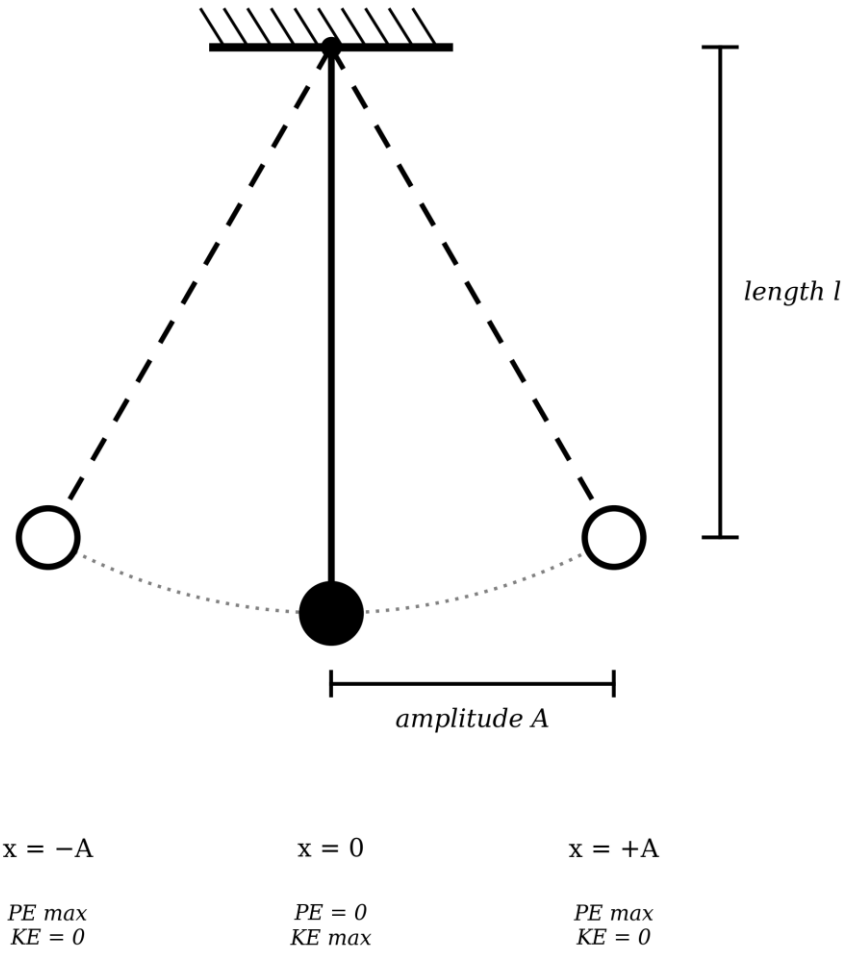


Figure: ***A simple pendulum at three positions of its oscillation.*** *At the extreme left (x = −A) and the extreme right (x = +A) the bob is momentarily at rest, so its kinetic energy is zero and its potential energy is maximum. At the equilibrium position (x = 0) the bob is moving fastest, so its kinetic energy is maximum and its potential energy is zero. The dotted curve marks the bob's path; the bracket on the right shows the pendulum length l; the bracket below shows the amplitude A.*

Period, frequency, and angular frequency

The **period**, written **T**, is *the time taken to complete one full oscillation:* that is, one full back-and-forth round trip that returns the body to the same position and the same direction of motion. Its SI unit is the second. A clock pendulum has a period of about 2s; a guitar string playing the note A above middle C has a period of about 2.27ms; the basilar membrane in the inner ear, when responding to a high-pitched whistle, may have a period as short as 0.05ms.

The **frequency**, written **f**, is *the number of complete oscillations per unit time.* It is the reciprocal of the period:

$$\mathbf{f = \frac{1}{T}}$$

Its SI unit is the **hertz** (1Hz = one oscillation per second), named after Heinrich Hertz, the physicist who first generated and detected radio waves in the laboratory about 140 years ago. A clock pendulum has a frequency of about 0.5Hz; a guitar string at the same A note has a frequency of 440Hz; the human ear can hear frequencies between roughly 20Hz and 20kHz, with the upper limit shrinking gently as we get older.

The **angular frequency**, written **ω** (the Greek letter omega), is the frequency expressed in radians per second rather than oscillations per second. Since one full oscillation corresponds to 2π radians, the conversions are:

$$\boldsymbol{\omega = 2\pi f} \quad \text{or} \quad \boldsymbol{\omega = \frac{2\pi}{T}}$$

Angular frequency may look like an unnecessary luxury *(why not stick with the friendly f in hertz?),* but every equation in this chapter will contain an ω rather than an f. The reason is geometric. We will derive the equations of oscillation by comparing them with uniform motion in a circle, and a particle moving in a circle naturally wants its speed described in radians per

second. Once you accept that one round trip is 2π radians of phase, the formulas come out clean and stay clean. We will let the swing of the pendulum itself decide which symbol to use, and ω will win nine times out of ten.

The energy picture

An oscillation is, at heart, a continual exchange between two forms of energy: **kinetic energy (KE**), *the energy of motion*, and **potential energy (PE)**, *the energy stored in the displacement away from equilibrium.* At the extreme positions, the body is momentarily at rest, so its KE is zero and all the energy of the system is locked up in PE. At the equilibrium position, the body is moving fastest, so its KE is maximum and its PE is zero. In between, the energy is shared in changing proportions, but if no energy is lost to friction or air resistance, the **total** energy KE + PE remains exactly **constant** throughout the motion.

This is the deepest pattern in the whole chapter, and it is worth pausing on. An oscillation is, in effect, a system that has too much of one kind of energy and is constantly trying to convert it into the other kind, and then back again. The conversion never stops. The total never changes. The figure earlier in this section showed the two extreme positions and the equilibrium position with the energy state of each marked beneath the bob. Look at it again now, and notice that the labels at the two extreme positions are the same: the system does not care which side of equilibrium it is on; it only cares how **far** from equilibrium it is. That symmetry is the oscillation speaking, and we will hear it speak many times before the chapter ends.

Free, forced, and damped oscillations

One last triple of vocabulary, then on to the equations. An oscillation is called **free** if *the only force acting on the body is the **restoring force** toward equilibrium*, with no external driver and no friction or air resistance to drain the energy away. A free oscillation has constant amplitude forever, and it occurs at one particular frequency, called the **natural frequency**, that is determined entirely by the properties of the system itself (the length of the pendulum, the stiffness of the spring, the density of the liquid). Free oscillations are an idealisation; in real life, every oscillation eventually loses energy to its surroundings.

An oscillation is called **damped** when *an external force, usually friction or air resistance, removes energy from the system on every cycle*. The amplitude of a damped oscillation decreases steadily over time, and the motion eventually dies out. Kipute's swing, left to itself after Kipanga walks away, will undergo damped oscillations: each cycle is a little smaller than the last, and after a few minutes the swing is back at rest. We will treat damping in detail later in the chapter, in a section devoted to it.

An oscillation is called **forced** when *an external periodic force is applied to keep it going*, as Kipanga's pushes did. A forced oscillation eventually settles into a steady state at the frequency of the driving force, not at the natural frequency of the system. When the driving frequency **matches** the natural frequency, however, something dramatic happens: the amplitude builds up to very large values, and the system is said to be in **resonance**. Resonance is the reason Kipanga's timed pushes worked when his random pushes did not. It is also the reason a tuned radio receives one station and not another, the reason a wineglass shatters at the right soprano note, and the reason an opera-singing physicist must be careful with the chandelier. The full treatment of resonance closes the chapter, and that is also where Kipanga returns to the swing for the last time.

These three words, free, damped, and forced, will return repeatedly in the chapters on waves, sound, and alternating current. They are the working vocabulary of every oscillating system in the syllabus, and the more comfortable you are with them now, the easier the rest of the book becomes.

THE DEFINING CONDITION OF SIMPLE HARMONIC MOTION

We have the vocabulary; we may now ask the proper question. *What must a back-and-forth motion obey, exactly, to qualify as **Simple Harmonic Motion (SHM)**?* The answer is one of the cleanest in physics. It is also one of the most powerful, because it turns out to apply to so many systems that you start to suspect nature is showing off.

The two conditions

Imagine that you take any oscillating system you like (a swing, a spring with a mass on it, a guitar string, a child on a mattress, a charged particle at the centre of a crystal, the bob of a pendulum) and you displace it a small amount from its equilibrium position. Two things must be true about the resulting motion if it is going to be simple harmonic.

Condition 1: *The acceleration of the body must be **proportional to its displacement** from equilibrium.* If you displace the body twice as far, it must accelerate twice as fast. If you displace it three times as far, it must accelerate three times as fast. The relationship must be exactly linear.

Condition 2: *The acceleration must always be directed **toward the equilibrium position***. Whenever the body is to the right of equilibrium, the acceleration must point left; whenever it is to the left, the acceleration must point right. The *acceleration must oppose the displacement*, every time, without exception.

Both conditions can be wrapped into one elegant equation. Let **x** be the displacement of the body from equilibrium and let **a** be the acceleration. Then SHM occurs if, and only if,

$$\mathbf{a = -\omega^2 x}$$

Where **ω** is some positive constant whose value depends on the particular system (the mass and the stiffness of the spring, the length of the pendulum and the strength of gravity, the depth and density of the liquid in the U-tube, and so on). The proportionality between a and x is captured by the factor ω^2. The fact that the acceleration points back toward equilibrium is captured by the minus sign. The constant ω is the angular frequency of the oscillation that the equation produces, the same ω we met in the vocabulary section.

Why the minus sign matters

It would be tempting, on first sight, to read that minus sign as a small algebraic detail and to forget it after one or two pages. Resist that temptation. The minus sign in $a = -\omega^2 x$ is the entire reason the motion oscillates rather than explodes.

To see why, suppose for a moment that someone wrote the equation without the minus sign:

$$a = +\omega^2 x$$

What kind of motion would this describe? Begin the body at a small positive displacement, $x > 0$. The equation tells us its acceleration is also positive, $a > 0$. So the body accelerates further in the positive direction. Its displacement grows. But as the displacement grows, the acceleration grows along with it (since a is proportional to x). So the body accelerates more and more rapidly, and its displacement runs away from equilibrium without bound. The "equilibrium" is no longer an equilibrium at all; it is a point that the body flees from at increasing speed. There is no oscillation, no rhythm, no return: just an exponential explosion.

The minus sign reverses all of this. With $a = -\omega^2 x$, a positive displacement produces a *negative* acceleration: the body is pushed back toward equilibrium. As it accelerates back, it overshoots equilibrium (because by the time it arrives there it has built up a velocity), and finds itself on the other side at a negative displacement, where the equation now says a is positive. It gets pushed back the other way. The result is a graceful, indefinite oscillation about $x = 0$, with the body returning, again and again, to the equilibrium it cannot help loving.

Therefore, the minus sign is the difference between a stable equilibrium (a body sitting in a valley, which oscillates) and an unstable one (a body balanced on a hilltop, which falls off). All simple harmonic motion lives in valleys. The minus sign is the valley. The figure below shows the picture in its simplest form: a parabolic potential well, with the body at three positions and the restoring force at each one pointing back toward $x = 0$.

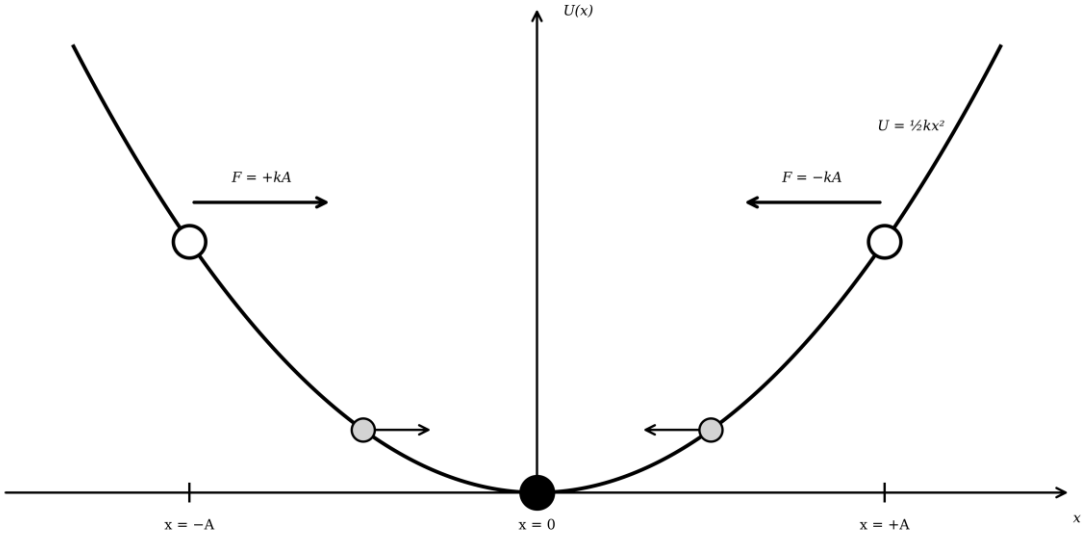


Figure: ***The parabolic potential well U(x) = ½kx² in which simple harmonic motion lives****. At every position other than equilibrium, the restoring force $F = -kx$ points back toward the bottom of the well. At the extremes ($x = \pm A$) the force has its maximum magnitude kA; at intermediate positions it is smaller; at the equilibrium $x = 0$ it is zero. This is the geometric meaning of the minus sign in $a = -\omega^2 x$: every disturbance is met with a push back home.*

The condition as a differential equation

Mathematically, the condition $a = -\omega^2 x$ is a **differential equation.** Acceleration is the second derivative of displacement with respect to time, so the same condition can be written in calculus form as:

$$\frac{d^2x}{dt^2} = -\omega^2 x$$

This is a request to find a function x(t) that, when differentiated twice with respect to time, returns a constant ($-\omega^2$) times itself. To search for such a function, we need an angular variable that grows linearly with time, so that the function repeats. Let us define this variable explicitly. As always, we use the Greek letter θ for an angle in radians. Here we set:

$$\theta = \omega t$$

The product ωt has units of radians, since ω is in radians per second and t is in seconds. So θ is a perfectly good angle, and it is the angle that grows linearly with time at angular speed ω. (In the next section we will see that θ has a beautiful geometric meaning: it is the angular position of a particle moving uniformly around a reference circle. For the moment we keep it as a definition.)

There are exactly two basic functions that satisfy the SHM equation: the sine and the cosine of θ. We check this by direct differentiation. Take $x = \sin\theta = \sin(\omega t)$. Differentiating once with respect to t (using the chain rule, since the inner function is ωt):

$$\frac{dx}{dt} = \frac{d}{dt}(\sin \omega t) = \omega \cos \omega t$$

Differentiating once more:

$$\frac{d^2x}{dt^2} = \frac{d}{dt}(\omega \cos \omega t) = -\omega^2 \sin \omega t = -\omega^2 x$$

So $x = \sin(\omega t)$ satisfies $d^2x/dt^2 = -\omega^2 x$. The same calculation for $x = \cos\theta = \cos(\omega t)$ gives:

$$\frac{dx}{dt} = -\omega \sin \omega t$$

$$\frac{d^2x}{dt^2} = -\omega^2 \cos \omega t = -\omega^2 x$$

So $x = \cos(\omega t)$ satisfies the SHM equation as well. Any linear combination of the two also passes the test (we shall demonstrate this in the second worked example below). The full family of solutions to the SHM differential equation is therefore

$$\mathbf{x(t) = a\sin \omega t + b\cos \omega t} \quad \text{(or equivalently } \mathbf{x(\theta) = a\sin \theta + b\cos \theta}\text{)}$$

Where a and b are constants determined by the initial position and initial velocity of the body. In the next section we will derive these solutions properly using a beautiful geometric trick (the reference circle), but for now we note only that they are sinusoidal: every simple harmonic motion is, when plotted against time, a sine curve. The two basic conditions on a and x have produced a single basic shape for the motion x(t).

We have the conditions, the equation, and the family of solutions. Now we put them to work. Three worked examples follow.

BINDER Example 1

Each of the following expressions describes a force F acting on a body of mass m moving along a coordinate axis (x, y, or angle θ). For each expression, decide whether the resulting motion would be simple harmonic, and justify your answer using the two conditions for SHM.

(a) $F = -2x^2$ (b) $F = -3y$ (c) $F = 10x$ (d) $F = -7\theta$

Solution

Apply Newton's second law in the form F = ma, so that

$$a = \frac{F}{m}$$

We then check the two conditions on the resulting acceleration: (1) is it proportional to the displacement? (2) is it directed back toward equilibrium (opposite in sign to the displacement)?

(a) $F = -2x^2$. Substituting into a = F/m:

$$a = \frac{-2x^2}{m}$$

The sign is negative when x is positive (so the acceleration points back toward equilibrium), but the dependence is on x^2, not on x. Condition 1 fails. Hence it is **not SHM.**

(b) F = -3y. Substituting:

$$a = \frac{-3y}{m} = -\left(\frac{3}{m}\right)y$$

Linear in the displacement and directed back toward equilibrium. Both conditions are met. Hence it is **SHM,** with $\omega^2 = 3/m$.

(c) F = 10x. Substituting:

$$a = \frac{10x}{m} = +\left(\frac{10}{m}\right)x$$

Linear in the displacement, but the sign is positive: the acceleration points in the **same** direction as the displacement, so the body is pushed **away** from equilibrium. Condition 2 fails. Hence it is **not SHM;** this is the exponential-explosion case discussed earlier.

(d) $F = -7\theta$. Substituting:

$$a = \frac{-7\theta}{m} = -\left(\frac{7}{m}\right)\theta$$

Linear in the angular displacement and directed back toward $\theta = 0$. Both conditions are met. Hence it is **Rotational SHM,** with $\omega^2 = 7/m$.

So only (b) and (d) describe simple harmonic motion.

Making Sense of the Answer: *The four expressions are designed to fail in different ways. (a) gets the direction right but the dependence wrong. (c) gets the dependence right but the direction wrong. Failing either condition kills SHM. The exercise is good practice in checking* ***both*** *conditions, not just the one that catches your eye first.*

Think Like a Physicist: *A common mistake is to look only at the minus sign and conclude that any force with a minus sign produces SHM. That is wrong: $F = -2x^2$ has a minus sign and is not SHM. The sign tells you the direction; the linearity in x tells you the proportionality. You need both. A second useful habit: once you have decided a force is SHM, read off ω^2 as the coefficient of the displacement (after dividing by m). The angular frequency is then immediate, and so is the period $T = 2\pi/\omega$.*

BINDER Example 2

A particle moves so that its displacement from equilibrium at time t is:

$$x(t) = a\sin\omega t + b\cos\omega t$$

Where a, b, and ω are constants. (i) Show that the motion is simple harmonic. (ii) Find the amplitude of the oscillation in terms of a and b. (iii) If a = 3cm and b = 4cm, find the amplitude.

Solution

(i) Differentiate x once with respect to t:

$$\frac{dx}{dt} = \frac{d}{dt}(a\sin\omega t + b\cos\omega t) = a\omega\cos\omega t - b\omega\sin\omega t$$

Differentiate once more:

$$\frac{d^2x}{dt^2} = \frac{d}{dt}(a\omega\cos\omega t - b\omega\sin\omega t) = -a\omega^2\sin\omega t - b\omega^2\cos\omega t$$

Factor $-\omega^2$ out of the right-hand side:

$$\frac{d^2x}{dt^2} = -\omega^2(a\sin\omega t + b\cos\omega t)$$

The bracket on the right is exactly x. Therefore:

$$\frac{d^2x}{dt^2} = -\omega^2 x,$$

which is the SHM condition. The motion is therefore simple harmonic with angular frequency ω.

(ii) The trick now is to combine the sine and cosine into a single sinusoid. Suppose we can write x as:

$$x = A\sin(\omega t + \varphi)$$

For some amplitude A and phase φ. Apply the sine-of-a-sum identity $\sin(\alpha + \beta) = \sin\alpha\cos\beta + \cos\alpha\sin\beta$ with $\alpha = \omega t$ and $\beta = \varphi$:

$$x = A(\sin\omega t\cos\varphi + \cos\omega t\sin\varphi) = (A\cos\varphi)\sin\omega t + (A\sin\varphi)\cos\omega t$$

Comparing this with the original $x = a\sin\omega t + b\cos\omega t$ term by term gives two equations:

$$A\cos\varphi = a \quad \text{and} \quad A\sin\varphi = b$$

Square each equation:

$$A^2\cos^2\varphi = a^2 \quad \text{and} \quad A^2\sin^2\varphi = b^2$$

Add the two squared equations:

$$A^2\cos^2\varphi + A^2\sin^2\varphi = a^2 + b^2$$
$$A^2(\cos^2\varphi + \sin^2\varphi) = a^2 + b^2$$

Use the Pythagorean identity $\cos^2\varphi + \sin^2\varphi = 1$ to simplify:

$$A^2 = a^2 + b^2$$
$$\mathbf{A = \sqrt{a^2 + b^2}}$$

The amplitude is therefore the square root of the sum of the squares of the two coefficients. (The phase angle is then found from $\tan\varphi = b/a$, although we did not need it here.)

(iii) Substitute a = 3cm and b = 4cm into $A = \sqrt{(a^2 + b^2)}$:

$$A = \sqrt{(3\text{cm})^2 + (4\text{cm})^2} = 5\text{cm}$$

The amplitude is 5cm.

Making Sense of the Answer: *We have shown two things at once. First, that any sum of a sine and a cosine of the same frequency is itself simple harmonic motion. Second, that the sum can always be rewritten as a single sinusoid with a single*

amplitude and a single phase. The amplitude of the combined motion is not (a + b), as a casual look might suggest, but $\sqrt{(a^2 + b^2)}$, the diagonal of a right-angled triangle whose two perpendicular legs are a and b. With a = 3cm and b = 4cm, that triangle is the famous 3-4-5 triangle from Form Two geometry, and the amplitude is a clean 5cm.

Think Like a Physicist: *The recipe in part (i) is one to remember: differentiate twice, factor out $-\omega^2$, notice that what remains in the bracket is x, then read off the SHM condition. The amplitude formula $\sqrt{(a^2 + b^2)}$ is the Pythagorean theorem in disguise, and once you have written it down once, you will recognise it everywhere: in the addition of two perpendicular SHM motions, in the addition of two waves of the same frequency, in the magnitude of an alternating-current phasor, in the resultant of two perpendicular electric fields. The pattern repeats. Almost every chapter from now until the end of the book will contain at least one moment where it appears.*

REAL Example 3

Read the following exchange. Then use the two conditions for simple harmonic motion to settle the argument once and for all.

Kipanga: *Sir, I have been thinking. Everything in the world is simple harmonic motion. The swing is SHM. The pendulum is SHM. The car suspension is SHM. Even the football that bounces on the playground when it has been kicked too softly to leave the ground, that is SHM, surely? It bounces up and down at a regular rate.*

Kipute: *Kipanga, no. A bouncing ball is not SHM.*

Kipanga: *But it is back-and-forth! It returns to the floor again and again! It has a period! What more do you want?*

Solution

Applying both conditions to the bouncing football.

Condition 1. *Is the acceleration proportional to displacement from some equilibrium position?* While the ball is in the air (which is where it spends nearly all of its motion), the only force on it is gravity, of magnitude mg directed downward. Its acceleration is therefore g downward, the *same value at every height*, regardless of how far above the ground it happens to be. Acceleration is constant, not proportional to displacement. The first condition for SHM fails immediately.

Condition 2. *Is the acceleration always directed toward an equilibrium position?* There is no equilibrium position in the air. The ball is in free fall, accelerating downward whether it is rising, falling, near the ground, or far from it. The only restoring force the ball ever feels is the brief contact force from the ground at the bottom of each bounce, lasting a small fraction of a millisecond. For almost the entire motion, the ball is in free fall with no restoring force at all. The second condition for SHM also fails.

The bouncing ball is therefore **nearly periodic, but not simple harmonic.** Kipanga has confused a much weaker condition (repetition at regular intervals) with the much stronger condition (a linear restoring acceleration toward equilibrium throughout the motion).

Making Sense of the Answer: *The temptation Kipanga fell into is a common one: to see back-and-forth motion at a regular rate and conclude it must be SHM. The conditions are stricter than that. SHM requires a restoring force that grows linearly with displacement, throughout the motion, with no jumps and no breaks. The bouncing ball spends most of its life in free fall, with no restoring force at all, so it is disqualified.*

Think Like a Physicist: *Most periodic motions in nature are not simple harmonic. The rotation of the Earth, the orbit of the moon, the bouncing of a ball, the beating of a heart, the firing of a neuron: all of these are periodic, none of them are SHM. SHM is the special case in which Hooke's law (or its analogue for angular motion, fluids, electric forces, or torsion) applies cleanly to the entire motion, not just part of it. That special case is, however, common enough (at the bottom of every smooth potential well in nature) to dominate physics in the same way that uniform circular motion dominates astronomy. Recognising what counts as SHM, and what does not, is the first true skill of this chapter.*

Three characteristics of simple harmonic motion

The equation $a = -\omega^2 x$ has three immediate consequences that we will use repeatedly throughout the chapter. We collect them here so you have them at your fingertips.

Characteristic 1: the period is independent of amplitude (the isochronous property)

The angular frequency ω depends only on the constants of the system (its mass, its stiffness, the local strength of gravity, the density of the liquid, and so on), not on how far you displaced the body to start the motion. The period $T = 2\pi/\omega$ is therefore also independent of amplitude. A pendulum displaced through 5° and the same pendulum displaced through 1° (assuming both angles are small) will swing at exactly the same period. This property is the reason the pendulum was a useful timekeeper for three centuries: as it loses energy to friction and its amplitude shrinks, the rate of swinging does not change, so the clock continues to tell time accurately even as it winds down.

Characteristic 2: the acceleration is maximum at the extremes and zero at equilibrium

From $a = -\omega^2 x$: when the displacement has its largest magnitude, $|x| = A$, the acceleration also has its largest magnitude, $|a| = \omega^2 A$. When the displacement is zero (the body is passing through equilibrium), the acceleration is zero. So the body is decelerating most strongly at the very moment it momentarily comes to rest at the extremes, and feels no net force at all as it speeds through the equilibrium position. This pattern of acceleration is the opposite of what intuition often suggests, and it is worth pausing on.

Characteristic 3: the total energy is conserved

The continual exchange of kinetic and potential energy that we described in the vocabulary section keeps KE + PE exactly constant throughout the oscillation, provided no energy is lost to friction or air resistance. We will derive the explicit expression for the total energy a few sections from now, but the principle is already in your hands.

These three characteristics, together with the equation $a = -\omega^2 x$ and the family of solutions $x(t) = a\sin\omega t + b\cos\omega t$, are the entire conceptual apparatus of simple harmonic motion. Everything from this point on in the chapter is, in one way or another, an application of these tools to a particular physical system. The next section turns the abstract solution into a concrete set of equations for displacement, velocity, and acceleration as functions of time, by means of the geometric trick mentioned earlier (the reference circle), which is, in our experience, the single most useful idea in the whole chapter.

THE EQUATIONS OF SIMPLE HARMONIC MOTION

We have shown that any simple harmonic motion can be written in the form $x(t) = a\sin\omega t + b\cos\omega t$, and that this expression can always be re-arranged as a single sinusoid $x(t) = A\sin(\omega t + \varphi)$. It is now time to give the individual equations of SHM in usable form. We will need three: one for the displacement x, one for the velocity v, and one for the acceleration a, each as a function of time. These equations, together with their phase relationships, are the calculation engine of the chapter. Almost every numerical question on SHM, from the simplest mass-spring exercise to the most ambitious damped-oscillator problem, can be solved by some combination of them.

Rather than solving the SHM differential equation by formal calculus, we will derive the equations using a beautiful geometric trick known as the reference circle. The trick goes back to Galileo, and it has the great merit of making the sinusoidal solution look obvious instead of mysterious. Once you have seen the trick once, you will not forget it.

The reference circle

Imagine a small particle P moving uniformly around a circle of radius **A**, in the counterclockwise sense, with constant angular speed **ω**. Place the centre of the circle at the origin O of an x-y coordinate plane. At time $t = 0$ let P be on the positive x-axis. After a time t has elapsed, P has swept out an angle $\theta = \omega t$ from the positive x-axis. Its coordinates at time t are therefore:

$$x_P = A\cos\theta = A\cos\omega t$$

$$y_P = A\sin\theta = A\sin\omega t$$

Now imagine that we drop a perpendicular from P down to the x-axis. The foot of that perpendicular, call it M, lies on the x-axis at the position $x = A\cos\omega t$. As P moves around the circle, M moves back and forth along the x-axis, between $x = +A$ (when P is on the right) and $x = -A$ (when P is on the left). The point M is the *projection* of P on the x-axis, and it executes simple harmonic motion. We will now demonstrate this directly as follows:

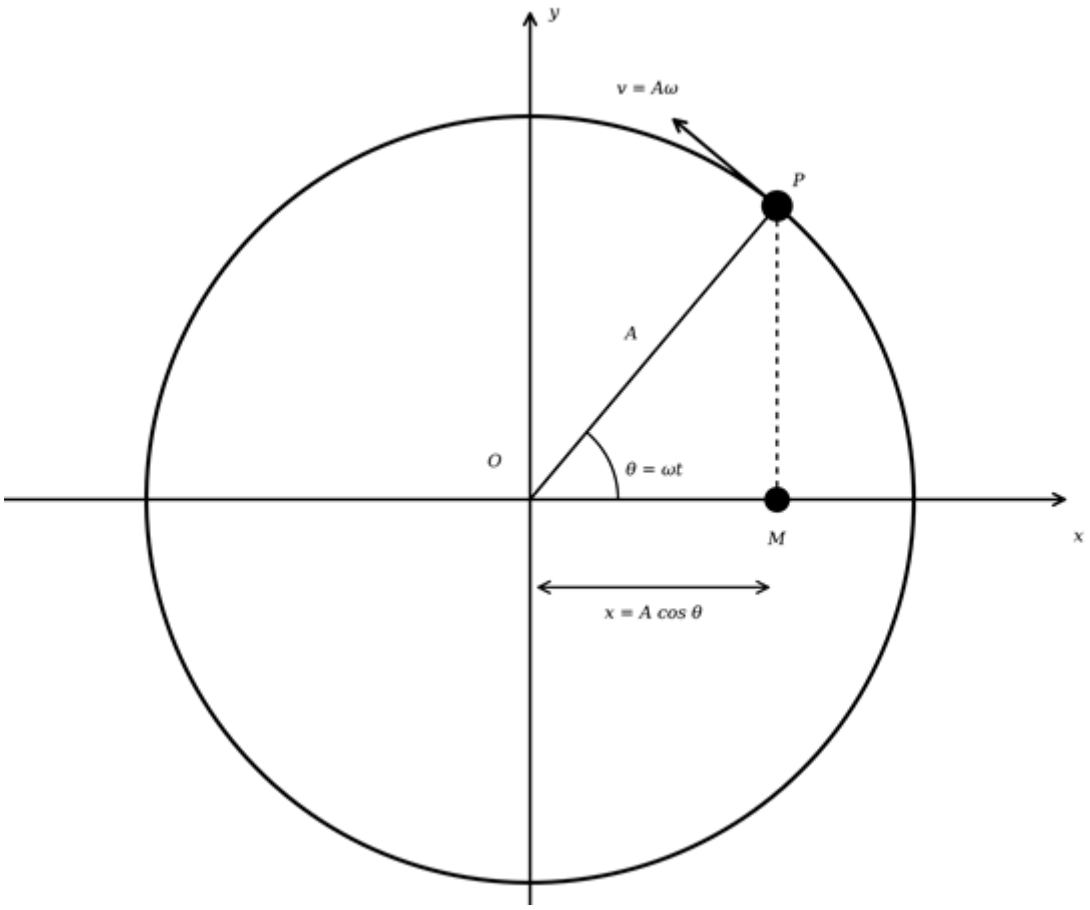


Figure: ***The reference circle for SHM****. A particle P moves uniformly around a circle of radius A at angular speed ω. Its angular position at time t is θ = ωt. The projection M of P onto the x-axis lies at x = Acos θ, and as P moves around the circle, M moves back and forth along the x-axis between x = +A and x = −A. The motion of M is simple harmonic.*

To prove that the motion of M is SHM, we need to show that its acceleration satisfies $a = -\omega^2 x$.

Differentiate the x-coordinate of M with respect to time once to obtain its velocity:

$$v = \frac{dx}{dt} = \frac{d}{dt}(A\cos \omega t) = -A\omega \sin \omega t$$

Differentiate once more to obtain the acceleration:

$$a = \frac{dv}{dt} = \frac{d}{dt}(-A\omega \sin \omega t) = -A\omega^2 \cos \omega t$$

Recognising that $A \cos \omega t = x$, the right-hand side becomes:

$$\mathbf{a = -\omega^2 x},$$

which is exactly the SHM condition. The projection M therefore obeys SHM with angular frequency ω, and the constants of the circle (its radius A and the angular speed ω) become directly the amplitude and the angular frequency of the SHM. The reference circle has done its work in two lines of calculus.

From this point on we have a complete physical picture: every simple harmonic motion can be thought of as the shadow of uniform circular motion on a diameter. The *radius* of the circle is the **amplitude**, the *angular speed around* the circle is the **angular frequency**, and one full *revolution* corresponds to **one full oscillation**. The **speed** of the circulating particle, which is constant and equal to **Aω**, becomes the **maximum speed** of the oscillation, attained as the projection M crosses the centre of the circle. We will use these relationships in the worked examples shortly.

Displacement, velocity, and acceleration

Two equivalent forms are commonly used for the displacement, depending on the initial condition. If the body starts at the equilibrium position (**x = 0 at t = 0**), then we choose the sine form:

$$\mathbf{x = A\sin \omega t}$$

If the body starts at the extreme position (**x = A at t = 0**), we choose the cosine form:

$$\mathbf{x = A\cos \omega t}$$

For the most general initial condition (**the body starts somewhere in between**), we use the form with an initial phase φ:

$$\mathbf{x = A\sin(\omega t + \varphi)}$$

From here on we will work with the sine form x = Asinωt unless a problem requires otherwise; the algebra is the same up to a re-labelling of t = 0.

Differentiate the displacement once with respect to time to obtain the **velocity:**

$$\mathbf{v} = \frac{dx}{dt} = \frac{d}{dt}(A\sin\omega t) = \mathbf{A\omega\cos\omega t}$$

The velocity is largest when cos ωt = ±1, which happens at t = 0, T/2, T, … (whenever the body is passing through equilibrium). The maximum speed is therefore

$$\mathbf{v_{max} = A\omega}$$

The velocity is zero when cos ωt = 0, which happens at t = T/4, 3T/4, … (whenever the body is at the extreme positions x = ±A). This matches the physical picture: the body must momentarily come to rest at the extremes before reversing direction.

It is often more useful to express the velocity in terms of the displacement x rather than the time t. We do this by eliminating t between the equations for x and v. Square both:

$$x^2 = A^2\sin^2\omega t$$
$$v^2 = A^2\omega^2\cos^2\omega t$$

Divide the second by ω², then add the two:

$$x^2 + \frac{v^2}{\omega^2} = A^2\sin^2\omega t + \frac{A^2\omega^2\cos^2\omega t}{\omega^2} = A^2\sin^2\omega t + A^2\cos^2\omega t$$

$$x^2 + \tfrac{v^2}{\omega^2} = A^2(\sin^2\omega t + \cos^2\omega t) = A^2 \text{ (using the identity } \sin^2\omega t + \cos^2\omega t = 1).$$

Solving for v gives the very useful form:

$$\mathbf{v = \pm\,\omega\sqrt{A^2 - x^2}}$$

This is the equation we will reach for whenever a question gives us the displacement and asks for the speed (or vice versa) without telling us the time. The ± sign reminds us that, at any given displacement other than the extremes, the body could be moving in either direction along the line: through equilibrium going right, or through equilibrium going left.

Differentiate the velocity once more to obtain the **acceleration:**

$$\mathbf{a} = \frac{dv}{dt} = \frac{d}{dt}(A\omega\cos\omega t) = \mathbf{-A\omega^2\sin\omega t}$$

Recognising that Asin ωt = x, we recover the SHM condition:

$$\mathbf{a = -\omega^2 x},$$

and we read off the maximum acceleration immediately. It occurs at the extremes (where |x| = A):

$$\mathbf{a_{max} = \omega^2 A}$$

The acceleration is zero at the equilibrium position, in agreement with the second of the three characteristics we listed in the previous section. The full set of equations is summarised below.

The equations of SHM (sine form):

- $x = A\sin\omega t$
- $v = A\omega\cos\omega t$ $\qquad v = \pm\,\omega\sqrt{A^2 - x^2}$
- $a = -A\omega^2\sin\omega t$ $\qquad a = -\omega^2 x$
- $v_{max} = A\omega$ $\qquad a_{max} = \omega^2 A$

Phase relationships between x, v, and a

The three quantities x, v, and a all oscillate sinusoidally at the same angular frequency ω, but they reach their peaks at different times. We can read the phase shifts directly from the equations.

Compare v = Aωcos ωt with x = Asin ωt. Using the identity cos θ = sin(θ + π/2), we can write the velocity as

$$v = A\omega\sin(\omega t + \pi/2)$$

So v has the same form as x but with the argument advanced by π/2 radians. We say that the velocity **leads** the displacement by a quarter cycle (π/2, or T/4). Whatever value of x occurs at time t, the same fractional value of v has already occurred a quarter period earlier.

The acceleration, on the other hand, is given by $a = -A\omega^2\sin\omega t$. Using -sin θ = sin(θ + π), we may write

$$a = A\omega^2\sin(\omega t + \pi)$$

So a has the same form as x but with the argument advanced by a half cycle (π, or T/2). The acceleration is **anti-phase** with the displacement: when x is at +A, a is at its most negative; when x is at −A, a is at its most positive. This is exactly what the equation $a = -\omega^2 x$ already tells us in compact form.

Putting the two phase shifts together: v leads x by T/4, and a leads v by T/4 (which makes a lead x by T/2, as the equation says). The following figure shows the three curves on the same time axis. The pattern repeats every period: x rises from zero,

v reaches its peak as x passes equilibrium, x reaches its peak as v passes zero, and a reaches its most negative value at the same instant.

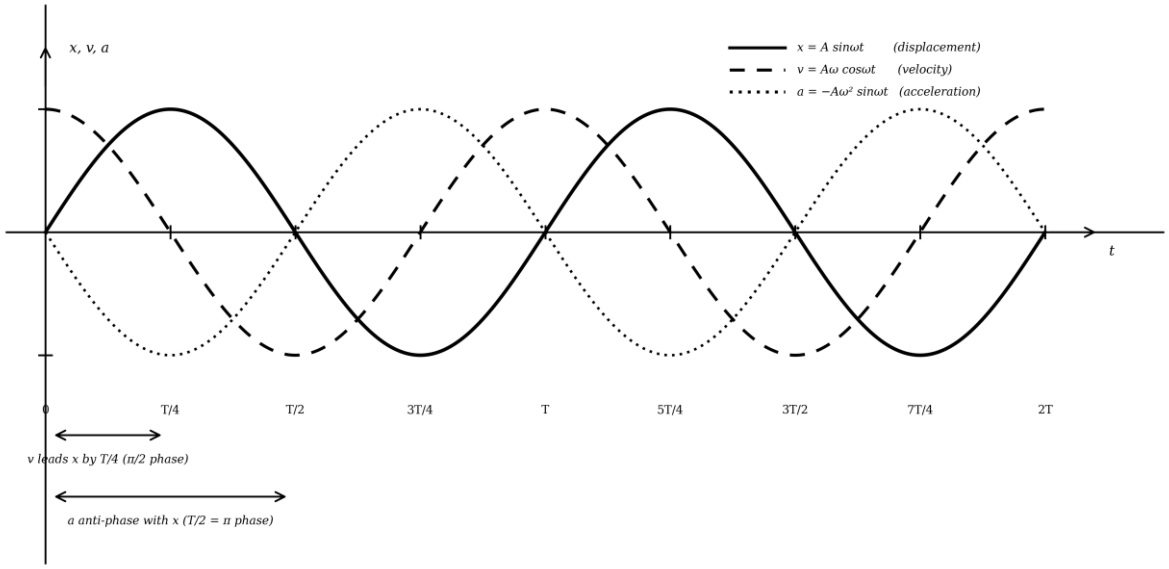


Figure: ***Displacement x = Asinωt (solid), velocity v = Aωcosωt (dashed), and acceleration a = −Aω² sinωt (dotted) plotted against time.*** *The three quantities oscillate at the same angular frequency ω but at different phases. The velocity leads the displacement by a quarter period (π/2), and the acceleration is anti-phase with the displacement (π, or half a period). The amplitudes have been normalised in the figure so that the phase relationships are clearly visible; in a real oscillation the actual peak values are A, Aω, and Aω² respectively.*

Five worked examples follow. They will work the equations of motion in every direction we are likely to be asked about in practice.

BINDER Example 4

A particle executes simple harmonic motion with a period T = 0.50s and amplitude A = 6.0cm. Calculate (a) the maximum speed of the particle, (b) the speed at the instant the particle is at half the amplitude.

Solution

First find the angular frequency ω from the period using $T = 2\pi/\omega$:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{0.50\text{s}} = 4\pi \text{ rad/s}$$

(a) Maximum speed occurs at x = 0 (the equilibrium position), and is given by $v_{max} = A\omega$:

$$v_{max} = (0.06\text{m}) \times (4\pi \text{ rad/s}) = 0.24\pi \text{ m/s} = 0.754 \text{ m/s}$$

The maximum speed is approximately 0.75m/s.

(b) Use the displacement-form $v = \pm \omega\sqrt{A^2 - x^2}$ with x = A/2 = 0.030m and A = 0.060m.
Substituting numerical values:

$$v = (4\pi \text{ rad/s}) \times \sqrt{(0.060\text{m})^2 - (0.030\text{m})^2} = 0.654\text{m/s}$$

The speed at x = A/2 is approximately 0.65m/s, taking the magnitude (the sign indicates direction along the line of motion).

Making Sense of the Answer: *At x = A/2, the speed is about 87% of the maximum speed. This is consistent with the fact that the particle loses very little speed as it moves from equilibrium out to halfway, and then rapidly decelerates over the second half of its journey to the extreme. The result also illustrates an important pattern: the speed of an SHM oscillator is not a linear function of position; it follows the elliptical relationship $v^2 = \omega^2(A^2 - x^2)$.*

Think Like a Physicist: *Always work in SI units (metres and seconds) before plugging numbers in, even when the question quotes centimetres and milliseconds. The angular frequency ω is naturally in rad/s, and the speed Aω comes out in m/s only if A is in metres. A common slip is to use A in cm and ω in rad/s, which gives a speed in cm/s; that is correct, but the unit must be reported faithfully. The shortest way to avoid confusion is to convert to SI at the start and to think only in SI thereafter.*

BINDER Example 5

A particle executing simple harmonic motion has an acceleration of 8.0m/s² at the instant when its displacement from the equilibrium position is 0.50m. Calculate the period of the motion.

Solution

Use the SHM relation $a = -\omega^2 x$, which gives the magnitude relation $|a| = \omega^2|x|$. Solve for ω^2:

$$\omega^2 = \frac{|a|}{|x|} = \frac{8.0\ \text{m/s}^2}{0.50\ \text{m}} = 16\ \text{rad}^2/\text{s}^2$$

Take the positive square root for ω:

$$\omega = \sqrt{16\ \text{rad}^2/\text{s}^2} = 4.0\ \text{rad/s}$$

The period follows from $T = 2\pi/\omega$:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{4.0\ \text{rad/s}} = \frac{\pi}{2}\ \text{s} = 1.57\text{s}$$

The period of the motion is approximately 1.57s.

Making Sense of the Answer: *An angular frequency of 4 rad/s is roughly that of a slowly-swinging child on a long swing, and a period of about 1.6s matches that physical picture nicely. The key step is the very compact use of $|a| = \omega^2|x|$, which gives ω directly from any single (a, x) measurement, without needing to know the amplitude or the time. This is one of the most useful relations in the chapter: every oscillator carries its angular frequency stamped on every (a, x) pair it visits.*

Think Like a Physicist: *Notice that the question gave us only one piece of data about each of two quantities (a and x at the same instant), and the SHM relation extracted ω from the pair. The amplitude A never entered the calculation, which makes physical sense, because ω depends only on the system, not on how violently it has been disturbed. If the same particle were released from a much larger amplitude, the relation between a and x at any instant would still give the same ω. The isochronous property is at work in the algebra.*

BINDER Example 6

The velocity of a particle executing simple harmonic motion is 16cm/s at a distance of 8cm from the mean position, and 8cm/s at a distance of 12cm from the mean position. Determine the amplitude of the motion.

Solution

Apply the displacement-form of the velocity relation, $v^2 = \omega^2(A^2 - x^2)$, at each of the two given instants. At the first instant, v = 16cm/s and x = 8cm:

$$(16\ \text{cm/s})^2 = \omega^2\left(A^2 - (8\ \text{cm})^2\right)$$

$$256\ \text{cm}^2/\text{s}^2 = \omega^2\left(A^2 - 64\ \text{cm}^2\right) \quad \ldots\ldots(\text{i})$$

At the second instant, v = 8cm/s and x = 12cm:

$$(8\ \text{cm/s})^2 = \omega^2\left(A^2 - (12\ \text{cm})^2\right)$$

$$64\ \text{cm}^2/\text{s}^2 = \omega^2\left(A^2 - 144\ \text{cm}^2\right) \quad \ldots\ldots(\text{ii})$$

Divide equation (i) by equation (ii) to eliminate ω^2:

$$\frac{256\ \text{cm}^2/\text{s}^2}{64\ \text{cm}^2/\text{s}^2} = \frac{A^2 - 64\ \text{cm}^2}{A^2 - 144\ \text{cm}^2}$$

$$A = 13.06\ \text{cm}$$

The amplitude of the motion is 13.06cm.

Making Sense of the Answer: *The two velocity readings, taken at two different displacements of the same oscillation, are enough to pin down the amplitude exactly, without needing to know either the period or the angular frequency. The trick is the $v^2 = \omega^2(A^2 - x^2)$ relation written twice and divided to cancel ω^2. The resulting amplitude (13cm) is larger than either of the two displacements (8cm and 12cm), as it must be, since both displacements are less than the amplitude.*

Think Like a Physicist: *This question pattern (two velocities at two displacements, find A) appears frequently in examination papers, and the recipe should become reflexive: write $v^2 = \omega^2(A^2 - x^2)$ twice, divide one equation by the other to kill ω^2, solve for A. Once you have A, you can also recover ω by substituting back into either equation, and from there get the period and the maximum speed if the question asks. The same trick is used elsewhere in physics whenever two readings of the same kind have been taken under different conditions, and you need to extract the system parameters.*

REAL Example 7

Read the dialogue below. Then use the equations for SHM to settle the question Kipanga raises.

Kipanga: *Kipute, this chapter is breaking my brain. I have just read that the acceleration of an SHM oscillator is **biggest** exactly when the oscillator is **not moving**. At the extreme positions, $x = \pm A$, the body is momentarily at rest, and yet that is where the acceleration is largest. I have been taught since Form Two that an object that is not moving has zero acceleration.*

Kipute: *No. Form Two never said that. Form Two said that an object in **equilibrium** has zero acceleration. An object that is **momentarily at rest** but **not in equilibrium** can have any acceleration at all.*

Kipanga: *I do not see the difference. The body is at rest. v = 0. So a = 0.*

Use the SHM equations to explain to Kipanga why his reasoning is wrong, and why the acceleration is in fact ***maximum*** at the extreme positions where the body is momentarily at rest.

Solution

Kipanga has confused two different physical situations. Velocity zero does not imply acceleration zero. Acceleration measures the rate at which velocity changes, not the value of velocity itself. We can have velocity zero at one instant and acceleration enormous at the same instant; the body is simply about to start moving rapidly in the opposite direction.

The SHM equations make this very precise. The acceleration is

$$a = -\omega^2 x$$

and the velocity is

$$v = \pm \omega\sqrt{A^2 - x^2}$$

These two equations say that a depends on x, while v depends on the difference between A^2 and x^2. At the extreme positions $x = \pm A$, the bracket $A^2 - x^2$ is zero, so v is zero. But $x = \pm A$ is also where $|x|$ is largest, so $|a|$ is at its maximum value, equal to $\omega^2 A$. Both the velocity-zero and the acceleration-maximum statements refer to the same instant in the same motion. There is no contradiction.

Making Sense of the Answer: *The key idea, both in this dialogue and in the SHM equations, is that velocity and acceleration are independent in the following sense: the value of one tells you nothing directly about the value of the other. They are related by the equation a = dv/dt, which is a* ***rate of change****, not a value. A high rate of change of velocity is compatible with the velocity itself being zero at that instant, just as a sharp turn in a road requires a large sideways acceleration even if the car is moving slowly.*

Think Like a Physicist: *This counter-intuition is one of the deepest sources of confusion in introductory mechanics, and it returns in many guises throughout physics: in projectile motion at the apex of the trajectory, in the spring-mass system at the turning points, in any oscillator at the extremes, and in alternating current circuits where the current can be zero at the very moment the rate of change of current is at its highest. The general lesson is that the value of a quantity and the value of its rate of change are different physical things; they do not have to vanish together, and they do not have to peak together. The pendulum swings most rapidly through the bottom and accelerates most strongly at the top. Once you have absorbed that pattern, the rest of oscillation physics becomes much easier.*

HOT Example 8

A particle oscillating with simple harmonic motion has, at the instant when it is 0.30m from its equilibrium position, a speed of 1.6m/s and an acceleration of magnitude 4.8m/s^2. Find (a) the angular frequency ω of the motion, (b) the amplitude A, (c) the period T, and (d) the maximum acceleration.

Solution

(a) From the SHM relation $|a| = \omega^2|x|$, we solve directly for ω:

$$\omega^2 = \frac{|a|}{|x|} = \frac{4.8\text{m/s}^2}{0.30\text{m}} = 16\ \text{rad}^2/\text{s}^2$$

$$\omega = \sqrt{16\text{rad}^2/\text{s}^2} = 4.0\ \text{rad/s}$$

(b) From the displacement-form of the velocity relation $v^2 = \omega^2(A^2 - x^2)$, solve for A^2:

$$A^2 = x^2 + \frac{v^2}{\omega^2}$$

$$A^2 = (0.30\text{m})^2 + \frac{(1.6\text{m/s})^2}{16\ \text{rad}^2/\text{s}^2}$$

$$A = \sqrt{0.250\text{m}^2} = 0.50\text{m}$$

(c) The period follows from $T = 2\pi/\omega$:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{4.0\text{rad/s}} = \frac{\pi}{2}\ \text{s} = 1.57\text{s}$$

(d) The maximum acceleration occurs at the extremes ($|x| = A$) and equals $a_{max} = \omega^2 A$:

$$a_{max} = (16\ \text{rad}^2/\text{s}^2) \times (0.50\text{m}) = 8.0\text{m/s}^2$$

Making Sense of the Answer: *All four quantities have come out to clean values. The amplitude A = 0.50m is comfortably larger than the given displacement x = 0.30m (as it must be, since the particle is on the way to or from one extreme but has not yet reached it). The maximum acceleration 8.0m/s^2 is larger than the given acceleration 4.8m/s^2 (as it must be, since the body was not at the extreme when measured), and indeed the ratio (8.0)/(4.8) = 5/3 equals exactly A/x = 0.50/0.30, which is the consistency check guaranteed by the SHM relation $a = -\omega^2 x$.*

Think Like a Physicist: *This is the most common multi-step SHM problem in the syllabus: given (v, a, x) at one instant, find (ω, A, T, a_{max}). The recipe is the same every time: (i) get ω from a/x, (ii) get A from $v^2 = \omega^2(A^2 - x^2)$, (iii) get T from $2\pi/\omega$,*

(iv) get a_{max} from $\omega^2 A$. Four steps, four well-defined quantities, and one consistency check (the ratio a_{max}/a should equal A/x). Once the recipe is in your fingertips you can solve any version of this problem in under five minutes.

With these equations in hand, every numerical question that the syllabus is likely to ask about a simple harmonic oscillator becomes a matter of picking the right formula and substituting carefully. The equations are the calculation engine of the rest of the chapter, and we will reach for them often. But the equations are silent about a feature of the motion that we promised, in the vocabulary section, to explore: the way the oscillator passes its energy back and forth between the moving body and the spring (or pendulum, or column of liquid, or whatever else is providing the restoring force). The energy is doing something specific at every position and at every instant, and the next section makes that something explicit.

ENERGY IN SIMPLE HARMONIC MOTION

An SHM oscillator carries energy in two forms. There is the kinetic energy of the moving body itself, which is largest when the body sweeps through the equilibrium position at maximum speed. And there is the potential energy stored in the restoring agent (the spring, the bent guitar string, the displaced liquid in the U-tube, the lifted bob of the pendulum), which is largest when the body is at the extreme of its excursion and the restoring agent is most strongly deformed. As the body oscillates back and forth, energy flows continually between these two forms, but the **total** energy stored in the system, undamped, never changes. We make this statement quantitative in this section.

Potential energy

The potential energy of an SHM oscillator is the work that must be done against the restoring force to move the body from the equilibrium position out to a displacement **x**. The restoring force on the body at position x is $\mathbf{F = -kx}$ (where **k** is the *stiffness of the system*, with $\mathbf{k = m\omega^2}$ in the SHM language). The agent doing the work against this force must apply a force +kx in the direction of the displacement. The work done by this external agent over a small further displacement **dx** is:

$$dW = kx\,dx$$

The total work done in moving the body from equilibrium (x = 0) out to displacement x is found by integrating:

$$U(x) = \int_0^x kx'\,dx' = \frac{1}{2}kx^2$$

Where x' is a dummy integration variable. So the potential energy of an SHM oscillator at displacement x is

$$\mathbf{U = \frac{1}{2}kx^2}$$

Using $k = m\omega^2$ (which we will derive properly when we treat the mass-spring system in the next section), this may be written equivalently as:

$$\mathbf{U = \frac{1}{2}m\omega^2x^2}$$

The potential energy is therefore proportional to the **square** of the displacement. It is zero at the equilibrium position (x = 0) and reaches its maximum value at the extremes (x = ±A), where its value is:

$$\mathbf{U_{max} = \tfrac{1}{2}kA^2 = \tfrac{1}{2}m\omega^2A^2}$$

Plotted against x, the potential energy is a parabola opening upward, exactly the shape of the potential well we drew earlier (page 338). SHM lives in a parabolic well, and the potential energy **is** the well.

Kinetic energy

The kinetic energy of the body at displacement x is $K = \tfrac{1}{2}mv^2$, with v given by the displacement-form of the SHM velocity equation:

$$v = \pm\,\omega\sqrt{A^2 - x^2}$$

Squaring,

$$v^2 = \omega^2\left(A^2 - x^2\right)$$

And substituting into $K = \tfrac{1}{2}mv^2$,

$$\mathbf{K = \frac{1}{2}m\omega^2\left(A^2 - x^2\right)}$$

Or equivalently $\mathbf{K = \tfrac{1}{2}k(A^2 - x^2)}$, since $k = m\omega^2$. The kinetic energy is therefore zero at the extremes (where x = ±A and the body is momentarily at rest) and reaches its maximum value at the equilibrium position (x = 0):

$$\mathbf{K_{max} = \frac{1}{2}m\omega^2A^2 = \frac{1}{2}kA^2}$$

Plotted against x, the kinetic energy is a parabola opening downward, reaching K_{max} at x = 0 and falling to zero at x = ±A. The two parabolas (the upward U-curve and the downward K-curve) are mirror images of each other across a horizontal line whose value we are about to determine.

Total energy is conserved

Add U and K at any displacement x:

$$E = U + K = \frac{1}{2}m\omega^2x^2 + \frac{1}{2}m\omega^2\left(A^2 - x^2\right)$$

The terms in x^2 cancel cleanly,

$$\mathbf{E = \frac{1}{2}m\omega^2A^2}$$

leaving the total energy independent of x and independent of t. The total energy of an undamped SHM oscillator is therefore **constant**, equal to $\frac{1}{2}m\omega^2A^2 = \frac{1}{2}kA^2$ at every instant. This is the third characteristic of SHM that we listed earlier, now derived rather than asserted.

Notice that *the total energy is proportional to the **square** of the amplitude.* Doubling the amplitude does not double the energy of the oscillator; it increases the energy by a factor of four. Trebling the amplitude multiplies the energy by nine. This square-law dependence is universal across all kinds of oscillators (mass-spring, pendulum, U-tube, guitar string, alternating-current circuit, sound wave) and it is one of the most important relationships in this chapter. We will return to it when we discuss damping, where the rate at which energy is lost from an oscillator is shown to be proportional to A^2, and the amplitude therefore decays exponentially with time.

Energy graphs

Two graphs of the SHM energies are useful enough that examiners ask for them by name. The first plots PE, KE, and total E against the displacement x; the second plots them against time t.

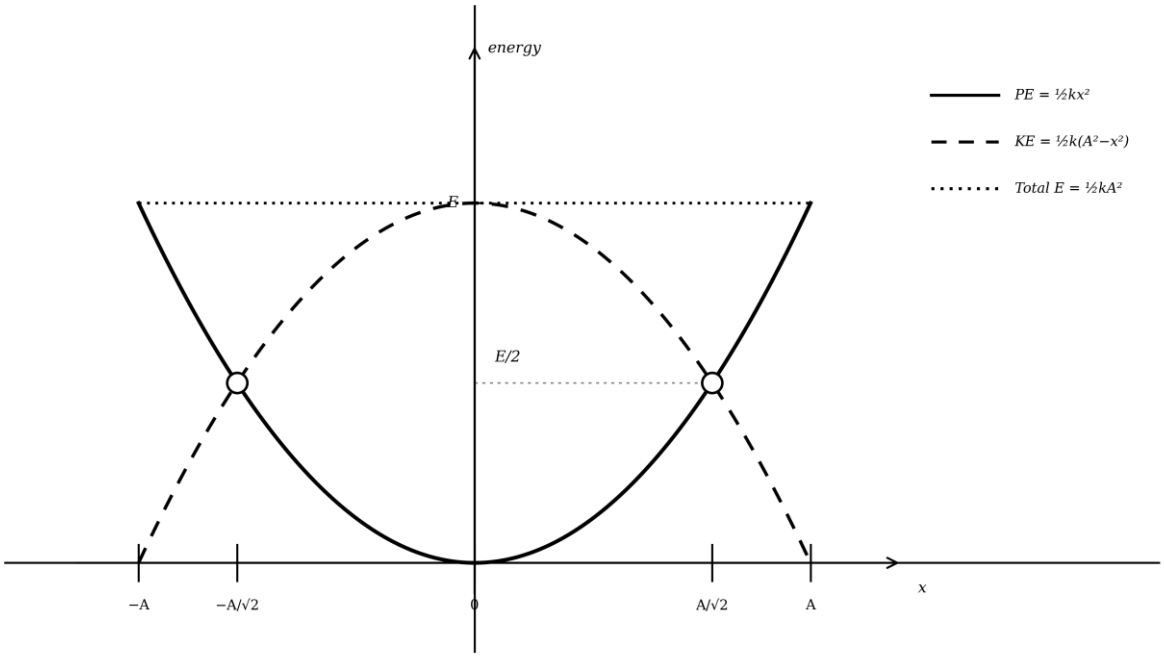


Figure: ***Energy of an SHM oscillator as a function of displacement x.*** *The potential energy $U = \frac{1}{2}kx^2$ (solid) is a parabola opening upward; the kinetic energy $K = \frac{1}{2}k(A^2 - x^2)$ (dashed) is a parabola opening downward; the total energy $E = \frac{1}{2}kA^2$ (dotted) is a horizontal line independent of x. The two curves cross at $x = \pm A/\sqrt{2}$, where $U = K = E/2$.*

From the energy-displacement graph, we can read directly an important position. The two parabolas cross when $U = K = \frac{1}{2}E$. Setting $U = \frac{1}{2}E$:

$$\frac{1}{2}kx^2 = \frac{1}{2}\left(\frac{1}{2}kA^2\right) = \frac{1}{4}kA^2$$

$$x^2 = \frac{A^2}{2}$$

$$x = \pm\frac{A}{\sqrt{2}} \approx \pm 0.707A$$

So the kinetic and potential energies are equal when the displacement is $(1/\sqrt{2}) \approx 0.707$ of the amplitude, or about 70.7% of the way out to the extreme. This is one of the most asked-about positions in the chapter, and it is worth memorising the result.

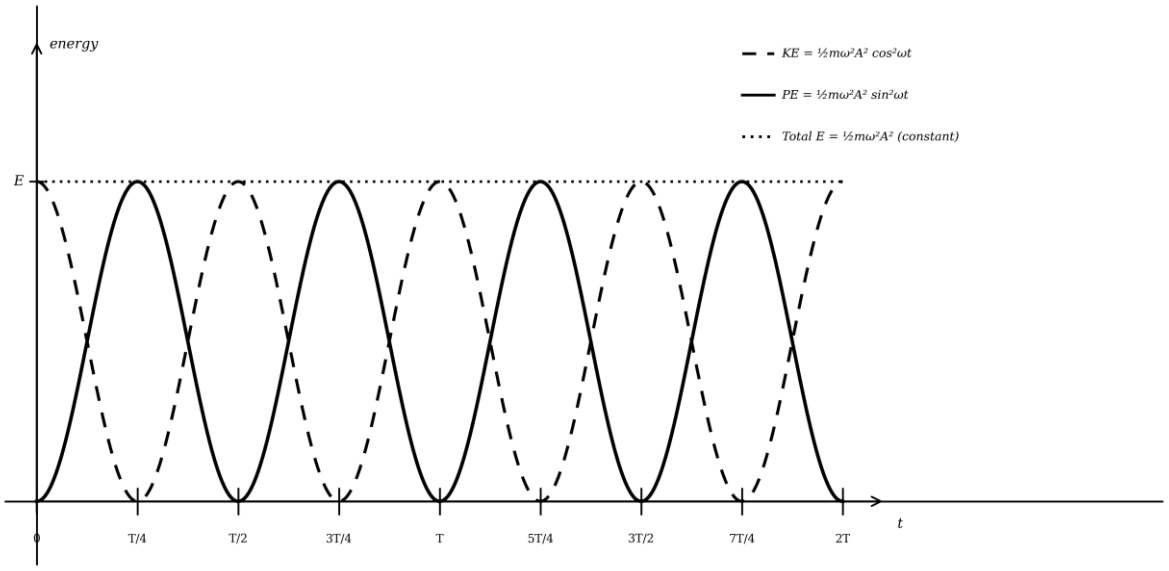


Figure: ***Energy of an SHM oscillator as a function of time t.*** *The kinetic energy K = ½mω²A²cos²ωt (dashed) and potential energy U = ½mω²A²sin²ωt (solid) both oscillate at twice the angular frequency of the displacement (period T/2 instead of T), and they are perfectly out of phase with each other: when one is maximum, the other is zero. The total energy E = ½mω²A² (dotted) is constant in time.*

The energy-time graph repays a moment of attention. Both KE and PE oscillate, but they oscillate at twice the frequency of the displacement itself (or equivalently with half the period). The reason is that energy depends on the **square** of position or velocity, and the square of a sinusoid completes a full oscillation in half the period of the original. So a body oscillating at frequency f exchanges its energy back and forth at frequency 2f. The total of course remains a flat line, since the sum of the two squared sinusoids is always equal to the constant E.

Four worked examples follow. They will look familiar in the same way that all the worked examples in this chapter look familiar; the trick is to know which formula they are quietly testing.

BINDER Example 9

A particle of mass 0.50kg vibrates with simple harmonic motion of period 1.0s. If its amplitude is 20cm, find its maximum kinetic energy.

Solution

First find the angular frequency ω from the period:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{1.0\text{s}} = 6.283 \text{ rad/s}$$

The maximum kinetic energy occurs at the equilibrium position (x = 0) and is given by $K_{max} = \frac{1}{2}m\omega^2A^2$. With m = 0.50kg, ω = 6.283rad/s, and A = 0.20m:

$$K_{max} = \frac{1}{2}(0.50\text{kg})(6.283\text{rad/s})^2(0.20\text{m})^2 = 0.395\text{J} \approx 0.40\text{J}$$

The maximum kinetic energy of the particle is approximately 0.40J.

Making Sense of the Answer: *An energy of 0.40 joules is what a half-kilogram brick gains by being lifted about 8cm against gravity, so the result is modest but tangible. Notice that the maximum kinetic energy depends on m, ω, and A together, and is sensitive to all three. If we doubled the amplitude we would quadruple K_{max}; if we halved the period (doubling ω), we would also quadruple it; doubling the mass would only double it.*

Think Like a Physicist: *The maximum kinetic energy equals the total energy of the oscillator. So the same formula ½mω²A² answers any question phrased as "maximum KE," "total energy," "maximum PE," or "energy stored in the oscillator." This is one of those formulas that quietly does the work of four when you know how to read it.*

BINDER Example 10

A particle is executing SHM of amplitude 5cm. At what distance from the mean position is the kinetic energy of the particle equal to its potential energy?

Solution

The kinetic and potential energies are equal when each equals half the total energy. Set the potential energy equal to half the total energy:

$$\frac{1}{2}kx^2 = \frac{1}{2}\left(\frac{1}{2}kA^2\right)$$

$$x^2 = \frac{A^2}{2}$$

$$x = \frac{A}{\sqrt{2}} = \frac{5\text{cm}}{\sqrt{2}} = 3.54\text{cm}$$

So the kinetic energy equals the potential energy at a displacement of about 3.54cm from the mean position, on either side.

Making Sense of the Answer: *The result is the same fraction of the amplitude ($1/\sqrt{2} \approx 0.707$) regardless of what the amplitude actually is. An oscillator with amplitude 5cm crosses its KE = PE point at 3.54cm; an oscillator with amplitude 1m crosses it at 70.7cm. The position is a property of the SHM, not of the particular oscillator. The point also marks the place where the body is moving at exactly ($1/\sqrt{2}$) of its maximum speed.*

Think Like a Physicist: *The KE = PE position is a hidden symmetry of the energy graph. Look at the energy-displacement figure: the two parabolas cross at exactly two points ($x = \pm A/\sqrt{2}$) and at no others. There is no oscillator anywhere in nature, of any size, made of any material, in which the kinetic and potential energies are equal at any displacement other than $\pm A/\sqrt{2}$. The result belongs to the geometry of the parabolas, not to the physics of any specific spring.*

REAL Example 11

Read the dialogue below. Then explain, using the energy formula for SHM, why Kipanga's prediction is wrong and what the correct answer is.

Kipanga: *Sir, last week you said that the total energy of an SHM oscillator is constant. So for the same spring with the same mass on it, if I displace it twice as far before letting go, the oscillator will have twice the energy when it starts swinging. Right?*

Kipute: *No. Four times the energy.*

Kipanga: *Four? But I only doubled the displacement. How does the energy quadruple from one extra factor of two?*

Solution

Kipanga has assumed that energy is proportional to amplitude. The SHM energy formula says otherwise. The total energy of an SHM oscillator is

$$E = \frac{1}{2}m\omega^2A^2,$$

and the amplitude appears as **a squared**, not as A. If we double the amplitude ($A \rightarrow 2A$), the new energy is

$$E_{new} = \frac{1}{2}m\omega^2(2A)^2 = 4 \times \left(\frac{1}{2}m\omega^2A^2\right) = 4E$$

So doubling the amplitude multiplies the energy by 4. If we trebled the amplitude, the energy would multiply by 9. If we increased it tenfold, the energy would multiply by 100. Energy depends on the **square** of amplitude.

Making Sense of the Answer: *The square-law dependence of energy on amplitude is responsible for many physical facts that, at first sight, look unrelated. A guitar string vibrating with twice the amplitude is four times as loud (in energy terms), not twice. A pendulum clock with twice the swing carries four times the energy through the air on each cycle, and loses energy correspondingly faster. A vibrating molecule with twice the amplitude stores four times the energy in the bond. The pattern is not a coincidence; it follows from the fact that potential energy in any Hooke-law system is quadratic in displacement.*

Think Like a Physicist: *Whenever you read a physical statement about an oscillator and find yourself thinking "linear in amplitude," stop and check. Almost no quantity associated with an SHM oscillator is linear in A. The maximum speed ($A\omega$) is linear in A; the maximum acceleration (ω^2A) is linear in A; almost everything else (energy, average power delivered, intensity, loudness, brightness, signal strength) is quadratic. The square is the rule, the linear the exception. Make a habit of asking which one applies.*

HOT Example 12

An object of mass 2kg executes SHM with a frequency of 2Hz and an amplitude of 2cm. Calculate (a) its maximum velocity and (b) its maximum potential energy. (c) At what position is the potential energy equal to three-quarters of the total energy?

Solution

First find ω from f:

$$\omega = 2\pi f = 2\pi \times 2.0\text{Hz} = 12.566\text{rad/s}$$

(a) Maximum velocity occurs at x = 0 and equals Aω. With A = 0.020m and ω = 12.566rad/s:

$$v_{max} = A\omega = (0.020\text{m}) \times (12.566\text{rad/s}) = 0.251\text{m/s}$$

The maximum velocity is approximately 0.25m/s.

(b) Maximum potential energy equals the total energy of the oscillator, ½mω²A². With m = 2kg:

$$U_{max} = \frac{1}{2}m\omega^2A^2$$

$$U_{max} = \frac{1}{2}(2\text{kg})(12.566\text{rad/s})^2(0.020\text{m})^2 = 0.0632\text{J} \approx 63.2\text{mJ}$$

The maximum potential energy is approximately 63.2mJ.

(c) Use the ratio U/E. From U = ½kx² and E = ½kA²,

$$\frac{U}{E} = \frac{x^2}{A^2}$$

Setting U/E = 3/4:

$$\frac{x^2}{A^2} = \frac{3}{4}$$

$$x = \frac{\sqrt{3}}{2}A = 0.866 \times 2.0\text{cm} = 1.73\text{cm}$$

The potential energy reaches three-quarters of the total energy when the displacement is approximately 1.73cm from the mean position, on either side.

Making Sense of the Answer: *The three results (v_{max} = 0.25m/s, U_{max} = 63mJ, x at 1.73cm) describe the same oscillator from three complementary angles. The maximum velocity is fixed by ω and A; the maximum potential energy is fixed by m, ω, and A; and the position where U = ¾E depends on A alone. None of the three depends on what kind of restoring agent is doing the work, and the same answers would apply to a 2kg mass on any spring tuned to ω = 12.6rad/s.*

Think Like a Physicist: *Part (c) used the very compact relation U/E = (x/A)². It is worth memorising in this form, because it lets you read off the displacement at which the energy is any fraction of the total. PE = ½E at x = A/√2; PE = ¾E at x = (√3/2)A; PE = ¼E at x = A/2; PE = E at x = A. Symmetric pairs of these positions divide every oscillation into energy-equal slices, and they are useful in problems on oscillator average behaviour and on the time spent in different parts of the cycle.*

We now have the equations of SHM, the energies of SHM, and the geometric picture of the reference circle. What we do not yet have is a single physical system whose oscillation we have analysed in full from first principles. We have spoken of springs and pendulums and U-tubes in passing, as examples of systems that obey a = −ω²x, but we have not derived the angular frequency of any of them from the system's own physics. That is the work of the next several sections.

We start with the simplest and most commercially useful of all oscillators, the mass on a helical spring, and we will see that the angular frequency falls out of Newton's second law and Hooke's law in two lines. The ω we have been carrying through every formula in this chapter will, for the first time, become a number we can compute from the mass of the body and the stiffness of the spring. From there the rest of the chapter follows: every other physical system we look at will be analysed by the same recipe, and each will give us its own ω.

THE MASS-SPRING SYSTEM

We have promised, since the very beginning of this chapter, that the angular frequency ω that appears in every SHM equation is determined entirely by the properties of the system itself, and is not a free parameter to be fitted from one problem to the next. We have also promised that for the simplest oscillating system in nature, a mass on a helical spring, the relationship comes out cleanly and quickly. The time has come to deliver on both promises. The mass-spring system is the workhorse of the chapter, the most exam-friendly oscillator in the syllabus, and the system most directly connected to everyday objects: car suspensions, pogo sticks, vibrating weighing scales, the up-and-down jiggle of a hanging chandelier when somebody slams a door. We start with the horizontal case, where the algebra is cleanest, then move to the vertical case, where gravity makes a cameo appearance and is then politely shown to make no difference, then finally to two springs combined either in series or in parallel.

Horizontal mass-spring system

Consider a mass m attached to one end of a horizontal helical spring of force constant k, with the other end of the spring fixed to a rigid wall. The mass rests on a frictionless surface, so the only horizontal force on it is the force from the spring. When the mass is displaced from its equilibrium position by a distance x (stretching or compressing the spring), Hooke's law gives the restoring force as:

$$F = -kx$$

Where the minus sign reflects that the force always points back toward equilibrium. Apply Newton's second law:

$$ma = F = -kx$$

Divide both sides by m:

$$a = -\frac{k}{m}x$$

Compare this with the SHM condition $a = -\omega^2 x$ term by term. The acceleration is proportional to the displacement and points back toward equilibrium, so the motion is simple harmonic. Reading off the angular frequency:

$$\omega^2 = \frac{k}{m}, \qquad \text{so} \qquad \omega = \sqrt{\frac{k}{m}}$$

and the period follows from $T = 2\pi/\omega$:

$$\mathbf{T = 2\pi\sqrt{\frac{m}{k}}}$$

This is the headline formula of the section, and it deserves a moment of contemplation. The period depends on the mass and on the stiffness of the spring, and on absolutely nothing else. It does not depend on the amplitude (because SHM is isochronous), it does not depend on the initial velocity, it does not depend on which direction you displaced the mass first, and it does not depend on gravity (because we are on a horizontal surface and gravity is balanced by the normal reaction of the table). A heavier mass means a longer period; a stiffer spring means a shorter period; both dependencies are square-root, so quadrupling either changes the period by a factor of two.

Vertical mass-spring system

Now hang the same spring vertically from the ceiling and attach the mass at the bottom. Gravity now pulls the mass downward with a steady force mg, in addition to the spring tension. At first sight you might worry that this changes the period: the mass has more force on it, after all. The first sight is wrong. The figure below shows what happens.

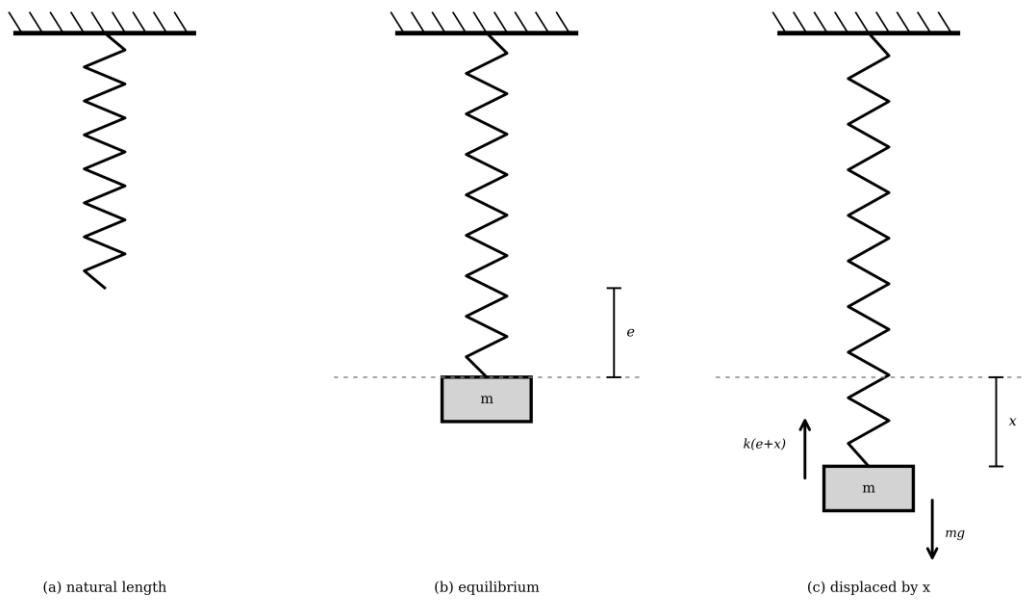


Figure: ***The vertical mass-spring system in three states.*** *(a) The spring at its natural length, no mass attached. (b) The mass attached and the system in equilibrium: the spring has stretched by a static extension e, and the upward spring tension ke balances the downward weight mg. (c) The mass pulled further down by an additional displacement x and released: the spring tension is now k(e + x), and the net downward force on the mass is mg − k(e + x) = −kx. The motion that follows is simple harmonic with the same period as the horizontal case.*

When the mass is in equilibrium under gravity alone (figure (b)), the spring is stretched by some static extension e. At equilibrium the upward spring tension ke balances the downward weight mg:

$$ke = mg$$

Now displace the mass by an additional distance x below this equilibrium position (figure (c)). The spring is now stretched by a total amount (e + x), so the upward spring tension is k(e + x). The downward weight is still mg. The net downward force on the mass is

$$F_{net} = mg - k(e + x) = (mg - ke) - kx = 0 - kx = -kx$$

since mg = ke at equilibrium. So the net force on the mass, measured from the equilibrium position, is just −kx, exactly the same as the horizontal case. Apply Newton's second law:

$$ma = -kx, \qquad \text{so} \qquad a = -\frac{k}{m}x$$

and the period is again:

$$\mathbf{T = 2\pi\sqrt{\frac{m}{k}}}$$

Gravity has not changed the period at all. What it changed was the equilibrium position: the mass now oscillates about the position it would settle at when left alone, not about the natural length of the spring. Once we measure displacement from the new equilibrium, the equation of motion is identical to the horizontal case. This is one of the most pleasant cancellations in physics, and it is worth pausing on. Hang any mass from any spring, in a uniform gravitational field of any strength, and the period of small oscillations is the same as it would be if you took the whole apparatus out into deep space and oscillated it horizontally. Gravity shifts the equilibrium; it does not change the rhythm.

There is one further useful relation that comes free with the vertical case. From mg = ke we have k = mg/e, so the period can be written entirely in terms of the static extension e and the local strength of gravity g:

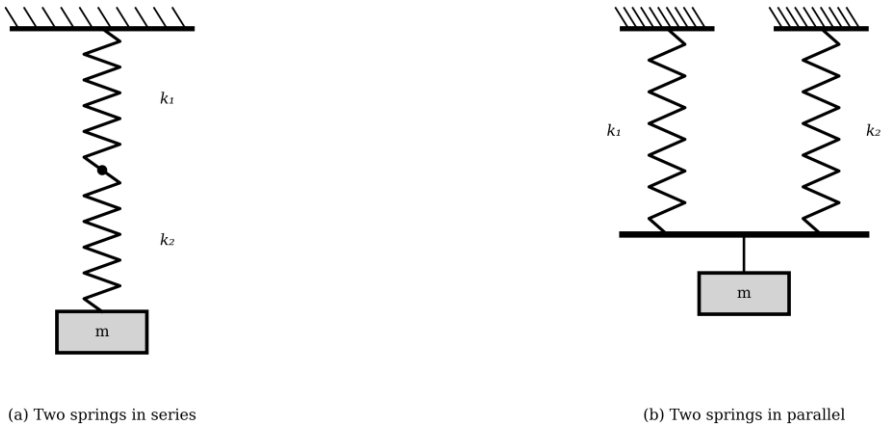
$$\mathbf{T = 2\pi\sqrt{\frac{e}{g}}}$$

This form is particularly useful in the laboratory, because the static extension e is much easier to measure than the spring constant k itself. Hang the mass on the spring, measure how far it descends, plug into the formula. A spring extended by 25mm in static equilibrium will give an oscillation period of about $2\pi\sqrt{(0.025/9.81)} \approx 0.32$s, regardless of what the spring is made of and regardless of what the mass actually is.

Recall also the chapter on elasticity, where we treated a 2kg mass hanging from a steel wire and asked for the period of small vertical oscillations. The answer came out at about 34ms. The same machinery is at work there: the steel wire behaves as a spring with k = YA/L (where Y is Young's modulus, A is the cross-sectional area, and L is the original length of the wire), and $T = 2\pi\sqrt{(m/k)}$ gives the very short period that surprised us at the time. The physical-world distinction between a spring and a wire is largely cosmetic; both store potential energy as $\frac{1}{2}kx^2$ and both produce SHM by the same mechanism.

Two springs combined: series and parallel

Many real systems use more than one spring. A car body sits on four springs (one above each wheel) and a diving board is essentially two long flexing springs working together. The two ways to combine springs, series and parallel, give different effective spring constants, and the period of oscillation depends on which is which. The figure below shows the two arrangements.



(a) Two springs in series (b) Two springs in parallel

Figure: ***Two ways to combine springs.*** *(a) Series: the two springs lie end to end, each carrying the full weight of the load. The total extension is the sum of the individual extensions. (b) Parallel: the two springs lie side by side, sharing the weight of the load. Both springs are stretched by the same amount.*

Series: Two springs of constants k_1 and k_2 are connected end to end, with a single mass hanging from the lower end. Both springs carry the same downward force F (the weight of the load is transmitted through the lower spring up into the upper spring without loss). The extension of spring 1 is $e_1 = F/k_1$ and the extension of spring 2 is $e_2 = F/k_2$. The total extension is

$$e = e_1 + e_2 = \frac{F}{k_1} + \frac{F}{k_2}$$

Defining the effective spring constant by $F = k_{eff} \times e$:

$$e = \frac{F}{k_{eff}} = \frac{F}{k_1} + \frac{F}{k_2}$$

$$\mathbf{\frac{1}{k_{eff}} = \frac{1}{k_1} + \frac{1}{k_2}}$$

For n identical springs of constant k in series:

$$\frac{1}{k_{eff}} = n \times \frac{1}{k}$$

$$\mathbf{k_{eff} = \frac{k}{n}}$$

The combined system is therefore **softer** than a single spring (smaller effective k), and the period of oscillation is **longer**, by a factor of √n.

Parallel: Two springs of constants k_1 and k_2 are arranged side by side, both attached to a rigid crossbar from which the load hangs. Both springs are stretched by the same amount e (because they are connected at both ends to rigid supports). The total upward force from the two springs is

$$F = F_1 + F_2 = k_1 e + k_2 e = (k_1 + k_2)e$$

So the effective spring constant is:

$$\mathbf{k_{eff} = k_1 + k_2}$$

For n identical springs of constant k in parallel, $\mathbf{k_{eff} = nk}$. The combined system is **stiffer** than a single spring (larger effective k), and the period of oscillation is **shorter**, by a factor of √n. A car on four springs in parallel oscillates twice as fast as it would on one spring of the same individual stiffness.

Five worked examples follow. The mass on the spring is the most heavily-tested oscillator in this chapter, so it is worth doing them slowly and not because they are exciting.

BINDER Example 13

A small mass of 0.2kg is attached to one end of a helical spring and produces an extension of 15mm. The mass is now pulled down 10mm and set into vertical oscillation of amplitude 10mm. Find: (a) the period of oscillation, (b) the maximum kinetic energy of the mass, and (c) the potential energy of the spring when the mass is 5mm below the centre of the oscillation.

Solution

(a) Use the static-extension form $T = 2\pi\sqrt{\frac{e}{g}}$ with e = 0.015m and g = 9.8m/s²:

$$T = 2\pi\sqrt{\frac{0.015\text{m}}{9.8\text{m/s}^2}} = 0.246\text{s}$$

The period of oscillation is approximately 0.25s.

(b) First find ω from the period:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{0.246\text{s}} = 25.57\text{rad/s}$$

The maximum kinetic energy occurs at the equilibrium position and equals $\frac{1}{2}m\omega^2A^2$.

With m = 0.2kg, ω = 25.57rad/s, and A = 0.010m:

$$K_{max} = \frac{1}{2}(0.2\text{kg})(25.57\text{rad/s})^2(0.010\text{m})^2$$

$$K_{max} = 6.54 \times 10^{-3}\text{J} \approx 6.5\text{mJ}$$

(c) First find the spring constant:

$$k = \frac{mg}{e} = \frac{(0.2\text{kg})(9.8\text{m/s}^2)}{0.015\text{m}} = 130.7\text{N/m}$$

Then $U = \frac{1}{2}kx^2$ at x = 0.005m:

$$U = \frac{1}{2}(130.7\text{N/m})(0.005\text{m})^2$$

$$U = 1.63 \times 10^{-3}\text{J} \approx 1.6\text{mJ}$$

Making Sense of the Answer: *All three answers are small in joules but quite manageable in millijoules, which is normal for laboratory-scale oscillators. Notice the relationship between the energies: the maximum kinetic energy (6.5mJ) is the total energy of the oscillation. The potential energy at x = A/2 = 5mm is 1.6mJ, which is exactly one-quarter of the total energy. That ratio (¼) arises because U/E = (x/A)², and (5/10)² = ¼.*

Think Like a Physicist: *The static-extension form $T = 2\pi\sqrt{(e/g)}$ is the most useful single formula in the section. It lets you compute the period of any vertical mass-spring system from a single ruler measurement of the static extension, without ever needing to know the spring constant or the mass separately.*

BINDER Example 14

When a metal cylinder of mass 0.2kg is attached to the lower end of a light helical spring, the upper end of which is fixed, the spring extends by 0.16m. The metal cylinder is then pulled down a further 0.08m and left free to oscillate vertically in simple harmonic motion. Calculate the period of vertical oscillations.

Solution

Use $T = 2\pi\sqrt{(e/g)}$ directly. The 0.08m "pulled down" is the amplitude of the oscillation, which (by the isochronous property) does not affect the period. Substituting e = 0.16m and g = 9.8m/s²:

$$T = 2\pi\sqrt{\frac{0.16\text{m}}{9.8\text{m/s}^2}} = 0.803\text{s}$$

The period of vertical oscillation is approximately 0.80s.

Making Sense of the Answer: *A period of about 0.8s for a 0.2kg cylinder on a spring that stretches 16cm under its weight is comfortable to imagine: roughly the rhythm of a slow, deliberate human heartbeat. The 0.08m amplitude information is deliberately included by the question setter as a distractor; the period of an undamped SHM oscillator is always independent of amplitude, and a careful reader will notice that no amplitude information appears in the formula $T = 2\pi\sqrt{(e/g)}$.*

Think Like a Physicist: *When you see numerical information in a question that appears to play no role in the formula, do not assume you have made an error. The information may be there to test whether you know it is irrelevant. The isochronous property is the most heavily tested fact about SHM, precisely because students who do not understand it are tempted to substitute the amplitude into formulas where it does not belong.*

BINDER Example 15

A 2.0kg ball is attached to a horizontal spring of negligible mass and force constant k = 50N/m on a frictionless table. The ball is displaced 0.20m from equilibrium and released. Find (a) the angular frequency, (b) the period, (c) the maximum speed, and (d) the maximum acceleration.

Solution

(a) Angular frequency from the formula $\omega = \sqrt{(k/m)}$:

$$\omega = \sqrt{\frac{k}{m}} = \sqrt{\frac{50\text{N/m}}{2.0\text{kg}}} = 5.0\text{rad/s}$$

(b) Period from $T = 2\pi/\omega$:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{5.0\text{rad/s}} = 1.26\text{s}$$

(c) Maximum speed at equilibrium, $v_{max} = A\omega$:

$$v_{max} = A\omega = (0.20\text{m}) \times (5.0\text{rad/s}) = 1.0\text{m/s}$$

(d) Maximum acceleration at the extremes, $a_{max} = \omega^2 A$:

$$a_{max} = \omega^2 A = (5.0\text{rad/s})^2 \times (0.20\text{m}) = 5.0\text{m/s}^2$$

Making Sense of the Answer: *A spring of k = 50N/m is fairly soft (about the stiffness of a typical retractable-pen spring), and 2kg is a moderately heavy mass for it. The resulting period of about 1.3s is leisurely, the maximum speed of 1m/s is gentle, and the maximum acceleration of 5m/s² is about half of g. All four quantities are in the comfortable middle range of laboratory oscillators, which is what a spring constant of 50N/m and a 2kg mass naturally produce.*

Think Like a Physicist: *Once you have ω from the system constants k and m, the rest of the standard SHM quantities follow from one formula each: period from $T = 2\pi/\omega$, maximum speed from $v_{max} = A\omega$, maximum acceleration from $a_{max} = \omega^2 A$, maximum KE from $K_{max} = \frac{1}{2}m\omega^2A^2$. The whole calculation collapses into four quick lines, and a question that asks for any of these is essentially testing whether you can navigate the formula tree from k and m at the root to the quantity asked for at the leaf.*

REAL Example 16

A car with a mass of 1300kg is constructed so that its frame is supported by four springs. Each spring has a force constant of 20000N/m. If two people riding in the car have a combined mass of 160kg, find the frequency of vibration of the car after it is driven over a pothole in the road.

Solution

The car body and its passengers oscillate together as a single mass M supported by four springs in parallel. Total oscillating mass:

$$M = 1300\text{kg} + 160\text{kg} = 1460\text{kg}$$

Springs in parallel combine by addition of spring constants. With four identical springs of k = 20000N/m:

$$k_{eff} = 4k = 4 \times 20000\text{N/m} = 80000\text{N/m}$$

Angular frequency:

$$\omega = \sqrt{\frac{k_{eff}}{M}} = \sqrt{\frac{80000\text{N/m}}{1460\text{kg}}} = 7.40\text{rad/s}$$

Frequency:

$$f = \frac{\omega}{2\pi} = \frac{7.40\text{rad/s}}{2\pi} = 1.18\text{Hz}$$

The car body oscillates at a frequency of about 1.2Hz after a pothole (which is roughly one full bounce per second).

Making Sense of the Answer: *A frequency of about 1Hz is the comfortable sweet spot for vehicle suspensions worldwide. Frequencies much higher than this feel jittery and make passengers uncomfortable on long journeys; much lower and the car rocks queasily after every bump. Engineers tune the spring constants and the damping (which we will treat later in this chapter) so that loaded family cars settle around this 1-1.5Hz range. The Tanzanian road context is real: anyone who has driven the Dar–Mbeya highway will testify that pothole-induced car oscillations are not just an academic example.*

Think Like a Physicist: *Notice how naturally the parallel-springs rule simplifies a real engineering question. Without it, you would have to track four separate forces from four separate springs, then add them. With it, the four springs behave as a single equivalent spring of constant 4k, and the period formula reads off in one step. The modular approach (replace any combination of springs by a single equivalent spring, then use $T = 2\pi\sqrt{(M/k_{eff})}$) scales to any number of springs in any combination of series and parallel, and is the practical way to design real oscillating systems.*

REAL Example 17

A delivery lorry plying the Dar es Salaam to Iringa highway carries sacks of maize. The driver notices that when the lorry is fully loaded with sacks, it bounces noticeably less violently after passing over a pothole than when it is running empty on the return trip. Using the mass-spring model of vehicle suspension, explain this observation. Address in particular: (i) how the period of bounce changes when the lorry is loaded, (ii) how the maximum speed of the bounce changes, and (iii) why the loaded ride feels gentler to the driver and to the cargo, even though more mass is being moved.

Solution

The body of the lorry rests on its leaf springs, and when the lorry hits a pothole the spring-mass system is set into vertical SHM about its equilibrium position. The springs together provide the same effective spring constant k_{eff} regardless of the load; what changes between loaded and empty is the total mass M of the body plus its cargo. Three SHM relations describe the bounce: the period $T = 2\pi\sqrt{\frac{M}{k_{eff}}}$, the maximum bounce speed $v_{max} = A\omega = \frac{2\pi A}{T}$, and the maximum bounce acceleration $a_{max} = \omega^2 A = \left(\frac{2\pi}{T}\right)^2 A$, where A is the amplitude (set by the depth of the pothole, which is the same on both trips).

(i) The period of bounce: Under the same set of springs, $T = 2\pi\sqrt{\frac{M}{k_{eff}}}$ varies as the square root of the total mass M. When the lorry is fully loaded, M is much larger than when it is empty. The period therefore increases in proportion to the square root of the increase in mass. As a result, the loaded lorry oscillates more slowly than the empty one.

(ii) The maximum speed of the bounce: With the same amplitude A but a larger period T, the maximum bounce speed $v_{max} = \frac{2\pi A}{T}$ is smaller when the lorry is loaded. Each bounce, although it moves the body through the same vertical distance, takes longer and therefore proceeds at a lower peak velocity. The cargo and the driver experience a gentler maximum speed of motion.

(iii) The acceleration the driver feels: The human body, and most cargo, senses jolts not as positions or speeds but as accelerations. The maximum bounce acceleration is $a_{max} = \left(\frac{2\pi}{T}\right)^2 A$, which falls as the inverse square of the period. So the loaded lorry has a longer period than the empty one, and thus its maximum acceleration is significantly reduced. This means the suspension responds more gently, so the driver experiences smaller jolts on the spine, the maize sacks bounce less violently, and the dashboard rattles much less.

This is why the loaded lorry feels more comfortable, even though it is moving more total mass.

The same principle is the reason that empty pickup trucks famously have a harsh ride while heavily-loaded lorries feel comparatively smooth to their drivers. The spring stiffness is the same; the mass is doing all the work.

Making Sense of the Answer: *The result is, at first sight, counter-intuitive. One might expect that a heavier load would strain the suspension more and produce a rougher ride. The physics says the opposite. More mass slows the oscillation and*

reduces both the bounce speed and (more strongly) the bounce acceleration. Comfort comes from gentle acceleration, and the simple inverse-square dependence of acceleration on period is what links a heavier load to a more comfortable ride.

Think Like a Physicist: *Real engineers exploit this systematically. Heavy goods vehicles are designed with springs stiff enough to support a fully-loaded mass, which means they ride harshly when empty (because then the effective mass is too small for the spring constant). Air suspensions on luxury buses go a step further: they vary the effective spring constant in real time as passengers board and disembark, keeping the period in the comfortable one-to-one-and-a-half hertz range regardless of how full the bus is. Behind every comfortable ride is a piece of SHM physics being applied deliberately, not accidentally.*

HOT Example 18

A tray of mass 4.0kg is supported by two identical springs. When the tray is displaced slightly and released, it executes simple harmonic motion with a period of 1.0s. (a) Find the spring constant of each spring. (b) A block of mass m is now placed on the tray; the period of oscillation increases to 2.0s. Find m.

Solution

(a) The two springs are in parallel (they support the tray side by side). The effective spring constant is $k_{eff} = 2k$. From $T = 2\pi\sqrt{\frac{M}{k_{eff}}}$:

$$T_1 = 2\pi\sqrt{\frac{M}{2k}}$$

Square both sides and solve for k:

$$T_1{}^2 = 4\pi^2\frac{M}{2k} = \frac{2\pi^2 M}{k}$$

$$k = \frac{2\pi^2 M}{T_1{}^2}$$

Substitute M = 4.0kg and $T_1 = 1.0s$:

$$k = \frac{2\pi^2 \times 4.0kg}{(1.0s)^2} = 79.0N/m$$

Each spring has a constant of approximately 79N/m.

(b) With the block of mass m on the tray, the total oscillating mass is $(M + m)$ and the new period is $T_2 = 2\pi\sqrt{\frac{M+m}{2k}}$. Squaring:

$$T_2^2 = \frac{4\pi^2(M+m)}{2k} = \frac{2\pi^2(M+m)}{k}$$

Divide this by the equation for T_1^2:

$$\frac{T_2^2}{T_1^2} = \frac{M+m}{M}$$

Solving for m:

$$m = M\left(\frac{T_2^2}{T_1^2} - 1\right)$$

$$m = (4.0kg)\left(\frac{(2.0s)^2}{(1.0s)^2} - 1\right) = (4.0kg)(4 - 1) = 12kg$$

The block has a mass of 12kg.

Making Sense of the Answer: *When the period doubled, the mass quadrupled. This is the inverse of the more familiar relation that the period varies as the square root of the mass. The recipe of dividing T_2^2 by T_1^2 to cancel the spring constant is one of the cleanest tricks in the chapter.*

Think Like a Physicist: *Whenever a problem gives you the same oscillator before and after some change (a mass added, a spring removed, the system taken from horizontal to vertical), look for a quantity that is unchanged by the change and use it to set up a ratio. Here the unchanged quantity is the spring constant 2k, and the ratio T_2^2/T_1^2 eliminated it automatically. Ratio methods of this kind almost always shorten a problem and reduce the number of arithmetic operations required.*

We have analysed the mass-spring system inside out: horizontal, vertical, in series, in parallel, with passengers, with extra blocks. Every variation reduced to one formula, $T = 2\pi\sqrt{(m/k)}$, and one geometric picture, the body sitting in a parabolic well of stiffness k. This is the recipe we will now apply to every other oscillating system in the chapter: identify the restoring force as a function of displacement, read off the effective stiffness, divide by the effective mass, take the square root and 2π, done. The next system on the list is the one Kipanga encountered first: the simple pendulum.

THE SIMPLE PENDULUM

The mass-spring system was kind to us because the spring did the heavy mathematical lifting. Hooke's law, F = -kx, is already in the SHM form before we even start. We did very little work to extract the period; we identified k and m, took a square root, and multiplied by 2π. The simple pendulum, is more interesting precisely because it has no spring. Its restoring force comes entirely from gravity, with no Hooke's law to lean on, and yet, after a careful look at the geometry, the SHM equation $a = -\omega^2 x$ reappears, with a new and beautiful meaning for ω.

The simple pendulum was Kipanga's very first encounter in this chapter. He tried, and failed, to push Kipute on the school swing. We now have the mathematical equipment to explain his failure and to dignify his eventual success. Behind every working clock pendulum, every grandfather clock that has measured weddings and burials for a hundred years, every metronome on every piano, the same equation lives. We are about to derive it.

What a simple pendulum is, and is not

A simple pendulum is an idealisation. It consists of a small, heavy bob, treated as a point mass, suspended from a fixed pivot by a long, light, inextensible string. The word "light" means we ignore the mass of the string entirely. The word "inextensible" means the string never stretches; its length is constant and equal to the distance from the pivot to the centre of the bob. The phrase "point mass" means the bob is small enough that we may treat it as concentrated at a single point.

No real pendulum is simple in this strict sense. The bob has finite size; the string has finite mass; and if the string is replaced by a metal rod, as in a clock pendulum, the rod is heavy enough that its own moment of inertia matters. Such a rod-and-disc arrangement is called a compound pendulum, and we shall treat it later in the chapter. The simple pendulum is the cleanest version of the device, the one that lets us isolate the gravity-and-length physics from any complication of mass distribution. Once we have the simple pendulum well in hand, the compound pendulum will be a small extension.

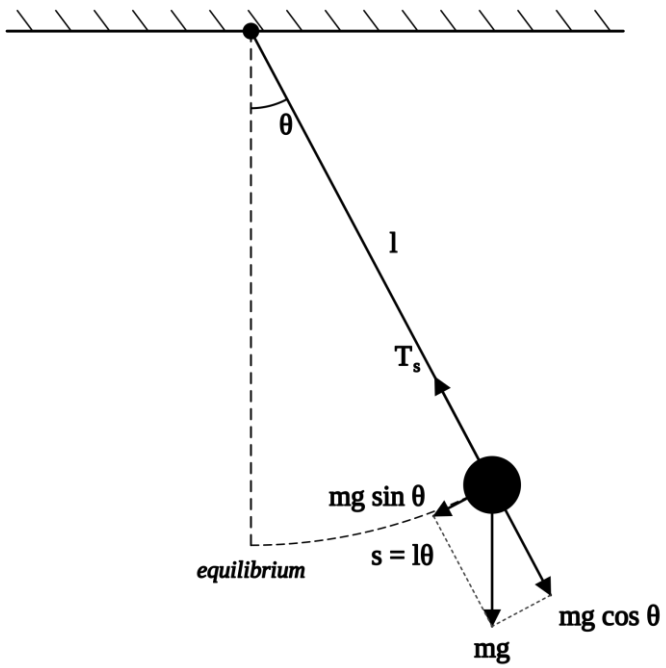


Figure: ***A simple pendulum displaced through a small angle θ from the vertical.*** *The bob hangs from a fixed pivot by a light inextensible string of length l. The arc length from the equilibrium position is s = lθ. Two forces act on the bob, the weight mg straight down and the tension T_s along the string. The weight resolves into a component mgcosθ along the string (balanced by the tension) and a component mgsinθ tangent to the arc, which is the restoring force always directed back toward equilibrium.*

From the swinging bob to the period formula

Place the pendulum at rest. The string hangs vertically, and the bob hangs at its lowest point. This is the equilibrium position, the only position at which the net force on the bob is zero. The two forces present, the weight mg straight down and the tension T_s straight up along the string, exactly cancel.

Now displace the bob through a small angle θ from vertical and release it. The bob is no longer at its lowest point; it stands at a **height l(1 - cos θ)** above equilibrium, swung along an arc of length **s = lθ** (with θ in radians). Two forces still act on the bob: the tension T_s along the string, now pointing up and back toward the pivot rather than vertically; and the weight mg, still pointing straight down.

To analyse the motion along the arc, we resolve the weight into two components. One component is along the string, of magnitude mgcos θ; the other is tangent to the arc, of magnitude mgsin θ. The component along the string is balanced by the tension. The component tangent to the arc is the restoring force. It always points back toward equilibrium, since when the bob is displaced to the right of vertical (positive θ), the tangential component points to the left, and the reverse holds on the other side. Calling the tangential force F_t:

$$\mathbf{F_t = -mgsin\ \theta}$$

The negative sign captures the restoring nature of the force, in the same way it did for the mass on a spring. Whenever θ is positive, F_t is negative; whenever θ is negative, F_t is positive. The force is always pulling the bob back home.

Newton's second law along the arc says that this tangential force equals mass times tangential acceleration. The tangential acceleration is the second derivative of arc length with respect to time, $a = \frac{d^2s}{dt^2}$, where s = lθ is the arc length. Therefore:

$$m\frac{d^2s}{dt^2} = -mgsin\ \theta$$

We have one equation. What we need now is to express the right-hand side in terms of the displacement s, so that we can compare with the SHM condition $a = -\omega^2 s$. Two substitutions do this. First, since s = lθ, we have θ = s/l. Second, for small angles measured in radians, sin θ ≈ θ. Putting both substitutions together:

$$sin\ \theta \approx \theta = \frac{s}{l}$$

So the equation of motion becomes:

$$m\frac{d^2s}{dt^2} \approx -mg\frac{s}{l}$$

The mass m cancels from both sides, a fact whose physical meaning we shall pause on shortly:

$$\frac{d^2s}{dt^2} = a = \approx -\frac{g}{l}s$$

This has exactly the form of the SHM condition $a = -\omega^2 s$, with:

$$\omega^2 = \frac{g}{l}$$

And therefore the period of the simple pendulum is:

$$\mathbf{T = \frac{2\pi}{\omega} = 2\pi\sqrt{\frac{l}{g}}}$$

This is the formula every Tanzanian advanced-level student is expected to write from memory. Let us look at it slowly. The period of a simple pendulum depends on only two things, its length l and the local strength of gravity g. It does not depend on the mass of the bob, nor on the amplitude of the swing (provided the small-angle approximation holds), nor on the colour, density, shape, or chemical composition of the bob. A pendulum of bronze, a pendulum of iron, and a pendulum of green mango, all swinging from strings of the same length on the same playground, have exactly the same period.

The cancellation of the mass is not a coincidence. It is the same cancellation that makes all bodies fall to the ground at the same rate: heavy bodies feel a larger gravitational pull, but they also have a larger inertia to be moved, and the two effects cancel exactly. The pendulum is therefore the cleanest possible test of the equality of gravitational and inertial mass.

The dependence on length is more useful in everyday practice. A pendulum twice as long as another has a period larger by a factor of $\sqrt{2} \approx 1.414$. To halve the period, you must shorten the length by a factor of four. To double the period, you must lengthen by a factor of four. This is why grandfather clocks have their pendulums about a metre long: that gives a period of just over two seconds, which the clock's escapement mechanism turns conveniently into the familiar tick-tock of one second per half-swing.

The dependence on g is much less noticeable in everyday life, since gravitational acceleration does not vary significantly from one place to another on Earth, but it becomes enormously important in scientific applications.

A pendulum is one of the most sensitive instruments ever invented for measuring g; a careful experiment can determine g to four significant figures with apparatus that costs less than a stove. We shall set up that experiment shortly.

When the motion is truly simple harmonic

The derivation above contained one approximation: sin θ ≈ θ for small θ in radians. For angles below about 10°, the small-angle approximation is correct, and the pendulum may be treated as performing simple harmonic motion. For 20° and above,

the approximation begins to fail noticeably, and the pendulum's period starts to depend on its amplitude; large swings take measurably longer than small ones, and the motion is no longer strictly SHM.

This is the precise answer to the recurring examination question, *"Why is the motion of a simple pendulum not strictly simple harmonic?"* The restoring force is $mg\sin\theta$, not $mg\theta$. The two are equal only in the limit of zero amplitude. Every real pendulum oscillation involves a small but non-zero amplitude, and is therefore only approximately simple harmonic. The approximation is excellent for the small angles that dominate practical use, but it must be acknowledged as an approximation, not as an exact identity.

We may now state the conditions on the simple pendulum precisely. The motion is approximately simple harmonic with period $T = 2\pi\sqrt{l/g}$ provided that:

1. the bob is treated as a point mass, with finite mass and zero size;
2. the string is treated as light (zero mass) and inextensible (constant length);
3. the angle of swing is small enough that $\sin\theta \approx \theta$; in practice, $\theta < 10°$ gives accuracy better than 0.5%;
4. the only forces on the bob are gravity and the tension in the string, with no air resistance and no friction at the pivot.

Whenever any of these conditions fails, the motion ceases to be strictly simple harmonic, although it may still be approximately periodic. Air resistance, in particular, causes the amplitude of the swings to decay over time, a topic we shall treat in detail in the section on damped oscillations later in this chapter.

The pendulum at zero gravity, and other places

Three thought experiments will fix the dependence of T on g in the imagination.

First, take a simple pendulum to the surface of the moon. The moon's surface gravity is about $g_{moon} \approx 1.62\text{m/s}^2$, about one-sixth of Earth's 9.81m/s^2. The same pendulum (same l) on the moon has a period:

$$\frac{T_{moon}}{T_{earth}} = \sqrt{\frac{g_{earth}}{g_{moon}}} = \sqrt{\frac{9.81\text{m/s}^2}{1.62\text{m/s}^2}} = 2.46$$

So a one-second-period pendulum on Earth becomes a 2.46-second-period pendulum on the moon. A clock built for Earth gravity would run grossly slow on the moon, and one would need to shorten its pendulum by a factor of about six to recover the same period.

Second, take the same pendulum to the top of Mount Kilimanjaro, 5895 metres above sea level. Local gravity decreases with altitude, since g varies as $1/r^2$, where r is the distance from the centre of the Earth. At the summit, g is reduced by about 0.18%. The fractional change in period is half of this, since $T \propto 1/\sqrt{g}$, so about 0.09%. A clock taken from Dar es Salaam to Uhuru Peak (highest point on Mount Kilimanjaro) loses roughly 80 seconds in a day. The effect is tiny but real, and it is routinely measured by precise gravity surveys.

Third, take the pendulum aboard the International Space Station, in orbit at 400km altitude. The Station and everything in it are in free fall around the Earth. The effective gravity inside is zero, which is what "weightlessness" in orbit actually means: not that gravity has vanished, but that everything is falling together. With $g_{eff} = 0$:

$$T = 2\pi\sqrt{\frac{l}{0}} = \infty$$

The pendulum simply does not oscillate. Pulled aside and released, it stays where it is, because in a frame falling freely with the bob there is no restoring force at all. Examiners enjoy this question, and the answer is unambiguous: a simple pendulum experiment cannot be done inside a satellite. *The pendulum needs a non-zero gravitational acceleration to provide its restoring force.* Take that away, and the device becomes inert.

Astronauts on the Space Station can, however, conduct mass-spring experiments perfectly well. The spring does not need gravity; it provides its own restoring force. This is one of the conceptual reasons mass-spring oscillators are studied alongside pendulums in the syllabus: they teach the same SHM equation but they fail in different environments, and a student who understands both can predict which experiment will work in which place.

Using the pendulum to measure g

Of all the experiments performed in school physics laboratories worldwide, the simple pendulum experiment for the determination of g is perhaps the most common, the most reliable, and the most rewarding. Its appeal lies in three properties at once: its apparatus is cheap (a length of string, a heavy bob, a stopwatch, a metre rule); its formula is clean ($T = 2\pi\sqrt{l/g}$ inverts to $g = 4\pi^2 l/T^2$); and the measurement uncertainty can be reduced indefinitely by patient repetition.

The basic procedure runs as follows:

1. Suspend a heavy bob (a brass cylinder, a steel ball with a hook, or a smooth stone with a string firmly tied) from a fixed pivot. A clamp on a retort stand is ideal. Measure the length l of the pendulum from the pivot to the centre of mass of the bob, using a metre rule.
2. Pull the bob aside through a small angle (less than 10° from vertical, to keep the small-angle approximation valid), and release it without giving it any initial sideways velocity.
3. Time the duration of **n** complete oscillations using a stopwatch. Choosing **n** large (say n = 20 or n = 50) reduces the influence of the human reaction time on each measurement, since the reaction time is fixed in absolute terms but the time for n oscillations grows linearly with n. The period is then:

$$\mathbf{T = \frac{total\ time}{n}}$$

4. Compute g from the formula:

$$\mathbf{g = \frac{4\pi^2 l}{T^2}}$$

A more refined version of the experiment varies the length l over five or six values, measures T at each, and plots T^2 (vertical axis) against l (horizontal axis). The graph should be a straight line through the origin, with a gradient equal to $4\pi^2/g$. The slope of this best-fit line gives g without ever needing to compute T at any single value; it averages over all the data and minimises systematic errors.

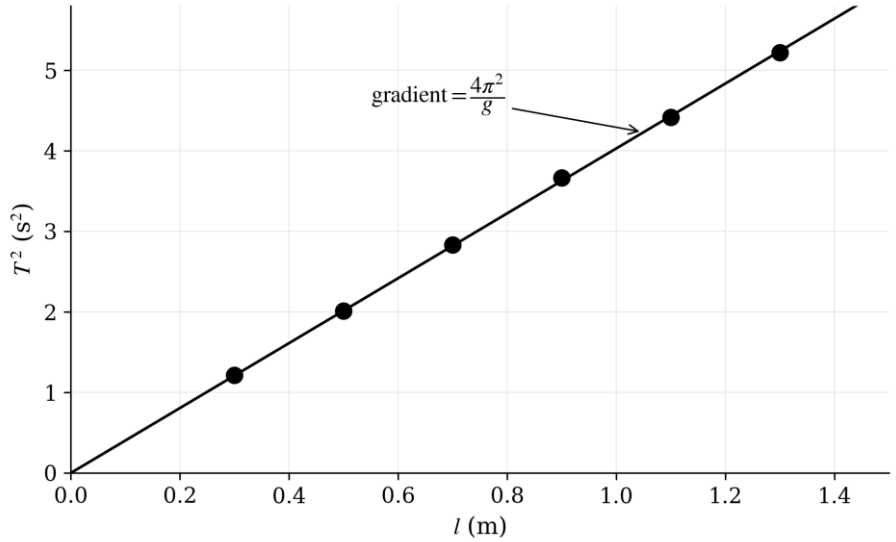


Figure: ***An idealised T² against l graph for a simple pendulum experiment.*** *Six measured points lie on a straight line through the origin. The gradient of the best-fit line equals 4π²/g, from which g may be computed without using any single (l, T) pair.*

The principal sources of uncertainty in this experiment are:

- **the measurement of l** (limited by the precision of the metre rule, typically ±0.001m, and by the difficulty of locating the centre of mass of the bob exactly), and
- **the timing of the oscillations** (limited by the human reaction time at the moments the stopwatch is started and stopped, typically ±0.1s on each press).

Propagating these uncertainties through $g = 4\pi^2 l/T^2$ gives:

$$\frac{\Delta g}{g} = \frac{\Delta l}{l} + 2\frac{\Delta T}{T}$$

The factor of 2 in front of $\Delta T/T$ comes from the fact that T appears squared in the denominator. Notice that the timing uncertainty contributes twice as heavily as the length uncertainty, fractionally, which is why students are urged to time many oscillations rather than few. Timing 100 oscillations instead of 10 reduces the fractional timing error by a factor of 10, and brings the dominant uncertainty down to the level of the length measurement.

Four worked examples now follow.

BINDER Example 19

In a school physics laboratory, a student timed 20 complete oscillations of a simple pendulum of length l = 100.0cm. The total time, measured by a stopwatch reading to the nearest 0.1s, was 40.0s. The metre rule used to measure the length reads to the nearest 0.1cm. (a) Calculate the value of g obtained in this experiment. (b) Estimate the maximum percentage error in the measured value of g.

Solution

(a) The period is the total time divided by the number of oscillations:

$$T = \frac{40.0\text{s}}{20} = 2.00\text{s}$$

The length is $l = 100.0\text{cm} = 1.000\text{m}$. Substituting into $g = 4\pi^2 l/T^2$:

$$g = \frac{4\pi^2(1.000\text{ m})}{(2.00\text{ s})^2}$$

$$g = 4\pi^2 \times 0.250\text{ m/s}^2 = 9.87\text{ m/s}^2$$

The obtained value of g is 9.87m/s^2.

(b) Propagating the errors:

Length uncertainty: $\Delta l = 0.1\text{cm} = 0.001\text{m}$, so $\Delta l/l = (0.001\text{ m})/(1.000\text{ m}) = 0.10\%$.

Timing uncertainty: the stopwatch reads to ±0.1s on each press, so the worst-case error in the total time is ±0.2s (one press at start and one at stop, each contributing ±0.1s). The error in $T = (\text{total time})/20$ is therefore $\pm 0.2\text{s}/20 = \pm 0.01\text{ s}$, so $\Delta T/T = (0.01\text{ s})/(2.00\text{ s}) = 0.50\%$.

Combined fractional error in g:

$$\frac{\Delta g}{g} = \frac{\Delta l}{l} + 2\frac{\Delta T}{T} = 0.10\% + 2 \times 0.50\% = 1.10\%$$

The maximum percentage error is 1.10%.

So $\Delta g \approx 1.10\% \times 9.87\text{ m/s}^2 \approx 0.11\text{ m/s}^2$. The result of the experiment, including its uncertainty, is therefore $g = (9.87 \pm 0.11)\text{ m/s}^2$.

Making Sense of the Answer: *The accepted value of g at sea level in Tanzania is approximately 9.81m/s^2. The experimentally determined 9.87m/s^2 lies within the calculated uncertainty of $\pm 0.11\text{m/s}^2$ of this accepted value, so the experiment is consistent with the textbook number, which is reassuring. Notice that the timing error dominates the uncertainty, contributing 1.0% out of the 1.1% total, while the length error contributes only 0.1%. This is the pattern in nearly every real pendulum experiment.*

Think Like a Physicist: *If you wanted to reduce the percentage error to below 0.5%, the cheapest improvement would not be a more precise metre rule, nor a more expensive stopwatch. It would simply be to time 100 oscillations instead of 20. The total reaction-time error remains ±0.2s, but it is now spread over a total time of about 200s instead of 40s, so $\Delta T/T$ falls by a factor of five, from 0.5% to 0.1%. The length error stays at 0.1%, and the combined fractional error becomes $0.1\% + 2 \times 0.1\% = 0.3\%$, comfortably under the half-percent target. This is the practical wisdom of the pendulum experiment: be patient with the timing, generous with the count, and the formula will do the rest.*

BINDER Example 20

Two simple pendulums of lengths $l_1 = 1.00\text{m}$ and $l_2 = 0.81\text{m}$ hang side by side from the same horizontal support and are set swinging at the same instant, in phase. After how many complete oscillations of the longer pendulum do they next swing exactly in phase again? Take $g = 9.8\text{m/s}^2$.

Solution

Let the period of pendulum 1 (longer) be T_1 and the period of pendulum 2 (shorter) be T_2.

From the simple-pendulum formula:

$$T_1 = 2\pi\sqrt{\frac{l_1}{g}}, \qquad T_2 = 2\pi\sqrt{\frac{l_2}{g}}$$

Each pendulum returns to the position and direction of motion it started with only after an integer number of complete oscillations. If pendulum 1 has performed n_1 complete oscillations, the time elapsed is $n_1 T_1$; if pendulum 2 has performed n_2 complete oscillations, the time elapsed is $n_2 T_2$. The two pendulums are observed at the same physical instant, so the two elapsed times must be the same number. That single requirement gives us the in-phase condition:

$$t = n_1 T_1 = n_2 T_2$$

Each side of this equation is the total elapsed time, written in two different ways. Both pendulums are simultaneously back to their starting phase only at instants when this equation is satisfied with positive integers n_1 and n_2. The next in-phase moment after release is the **smallest** such pair of integers.

Rearrange to isolate the ratio of integers:

$$\frac{n_1}{n_2} = \frac{T_2}{T_1} = \sqrt{\frac{l_2}{l_1}}$$

The square root in the last expression came from the period formulas: $T_2/T_1 = \sqrt{l_2/l_1}$, with the common factor $2\pi/\sqrt{g}$ cancelling.

We now substitute numbers. With $l_1 = 1.00\text{m}$ and $l_2 = 0.81\text{m}$:

$$\frac{n_1}{n_2} = \sqrt{\frac{0.81\text{m}}{1.00\text{m}}} = 0.9 = \frac{9}{10}$$

The smallest positive integers satisfying $n_1/n_2 = 9/10$ are $n_1 = 9$ and $n_2 = 10$. (The next in phase after this will occur at $n_1 = 2\times9=18$, $n_2 = 2\times10=20$.)

So to be in phase again the longer pendulum completes 9 oscillations (while the shorter completes 10).

Checking the answer:

The total elapsed time, in either form, is:

$$t = n_1 T_1 = 2\pi\, n_1 \sqrt{\frac{l_1}{g}}$$

Substituting:

$$t = 9 \times 2\pi \sqrt{\frac{1.00\text{m}}{9.8\text{m/s}^2}} = 18.06\text{s}$$

Or equivalently, $t = 10T_2 = 10 \times 1.806\text{s} = 18.06\text{s}$, as it must be.

Making Sense of the Answer: *The clean 9:10 ratio is no accident. The two lengths, 1.00m and 0.81m, were chosen deliberately so that $l_2/l_1 = 0.81 = (9/10)^2$. Taking the square root gives $T_2/T_1 = 9/10$, so the simplest integer match is $n_1 : n_2 = 9 : 10$. If the lengths had been chosen randomly, the integers could have been very large indeed; in the worst case, the lengths are incommensurable and the pendulums never return exactly to phase. The choice of lengths in this question is therefore a small kindness to the student, made obvious by the fact that $0.81 = 0.9^2$ is a perfect square in disguise. Notice also a methodological point: by working symbolically first and substituting numbers only at the end, we obtained the answer with one cleaner formula, $t = 2\pi n_1\sqrt{(l_1/g)}$, rather than two separate period calculations followed by a ratio. Wherever the algebra allows, this is the better order of operations.*

Think Like a Physicist: *The principle at work here is the simplest possible case of resonance, although that word will not officially appear in this chapter until the section on forced oscillations. Two oscillators with a simple integer ratio of their periods come into phase often, exchange energy efficiently if coupled, and produce all the phenomena of musical harmony. Two oscillators whose period ratio is irrational, the golden ratio for example, almost never come into phase, never exchange energy efficiently, and produce the perpetual unfinished sound of an out-of-tune instrument. The integer ratios of the pendulum lengths in this question are the same integer ratios that govern the consonant intervals of music. We shall meet them again, in different clothes, in the chapter on stationary waves.*

REAL Example 21

In the staffroom at Miono Secondary, there is a wall clock whose pendulum is a brass rod with a heavy disc at the end. The clock has hung on that wall for years, and the headmaster's secretary, who has the unenviable job of calling the start and end of every lesson, has noticed something peculiar. In the cool months between June and August, when the school sits at twenty degrees Celsius from morning to night, the clock keeps perfect time. In the hot months between February and April, however, when the temperature inside the building reaches thirty-five degrees Celsius and the brass pendulum visibly glistens, the clock loses about twenty seconds a day. By the end of a hot week she has had to march into the staffroom and reset it.

Kipute, who has been reading ahead, raises this with Mr. Akilikubwa one Friday afternoon.

Kipute: *"Sir, the wall clock is losing time again. The secretary is unhappy. Can you explain to her, in physics, why a hot brass pendulum is slower than a cool one?"*

Mr. Akilikubwa, who has been waiting since the chapter on thermal expansion to be asked exactly this question, smiles a small triumphant smile.

Mr. Akilikubwa: *"The pendulum is a brass rod. Brass, like all metals, expands when it is heated. From Chapter 2, the linear expansion of a rod under a temperature change ΔT is given by:"*

$$l = l_0(1 + \alpha\Delta T)$$

Mr. Akilikubwa: *"where α is the coefficient of linear expansion of brass, equal to about 19×10^{-6}/K, and ΔT is the temperature increase. Now, the period of a simple pendulum (and to a good approximation a rod-and-disc pendulum behaves like a simple pendulum of effective length close to the rod length) is:"*

$$T = 2\pi \sqrt{\frac{l}{g}}$$

Mr. Akilikubwa: *"Substitute the expanded length:"*

$$T = 2\pi\sqrt{\frac{l_0(1+\alpha\,\Delta\,T)}{g}} = T_0\sqrt{1+\alpha\Delta T}$$

Mr. Akilikubwa: *"For the small fractional changes typical of thermal expansion, we may expand the square root to first order, using* $(1+x)^{\frac{1}{2}}=1+\frac{x}{2}$ *for small x:"*

$$T \approx T_0\left(1+\frac{1}{2}\alpha\Delta T\right)$$

Mr. Akilikubwa: *"So the fractional change in period is half the fractional change in length. Plugging in numbers, a temperature increase of* $\Delta T = 15K$ *from 20°C to 35°C gives a fractional length change:"*

$$\frac{\Delta l}{l_0} = \alpha\Delta T = (19\times10^{-6}/K)(15K) = 2.85\times10^{-4}$$

Mr. Akilikubwa: *"and therefore a fractional period change of half this, about* 1.43×10^{-4}*. In a day of 86400s, the clock will be slow by:"*

$$(86\,400s)(1.43\times10^{-4}) \approx 12.4s \text{ per day}$$

Mr. Akilikubwa: *"Twelve seconds a day, in round figures, in agreement with the secretary's complaint of about twenty seconds, a number subject to the rounding errors of human observation. The clock loses time in summer because its pendulum rod is thermally longer in summer."*

Kipute: *"And what does the clockmaker do about it?"*

Mr. Akilikubwa: *"Two things, traditionally. The simple solution is to leave a small adjustment screw at the bottom of the rod that lets the position of the disc be moved up or down by a fraction of a millimetre. Moving the disc up shortens the effective pendulum length and speeds the clock; moving it down slows it. A patient secretary can compensate the seasonal change by hand."*

Mr. Akilikubwa*: "The clever solution, invented by* ***John Harrison*** *in eighteenth-century England, uses two metals with different expansion coefficients arranged so that their expansions cancel. As the brass rod gets longer in summer, a steel rod, which expands less, on the other side of the system gets longer too, but by less, and the lever arrangement keeps the effective pendulum length constant. Such a pendulum is called a temperature-compensated pendulum, and a well-built one keeps time to within a second a week regardless of the season."*

Kipanga, who has just sauntered in for water and stayed for the conversation, looks up.

Kipanga: *"So if I want the clock to be perfect, I should hire* ***Harrison****?"*

Mr. Akilikubwa: *"Harrison died in 1776. But you may hire any honest watchmaker on Samora Avenue, and they will do the same job today, for less."*

Making Sense of the Answer: *The whole calculation rests on three ideas you already have: thermal linear expansion (Chapter 2), the pendulum period formula (this section), and the binomial approximation (BAM/Advanced mathematics). The numerical result, about twelve seconds a day for a fifteen-degree temperature change, matches what is observed in school clocks across Tanzania every dry season, and is one of the most direct everyday consequences of the temperature-dependent length of solids. The same physics is the reason a railway track buckles in the heat and a power line sags in summer; it just appears, in this case, in a slightly more elegant disguise.*

Think Like a Physicist: *This problem is a small example of a very large theme in physics: the physics of small effects. The fractional period change of* 1.4×10^{-4} *is impossible to notice in any single oscillation; you would not see it with a stopwatch. But twelve seconds compounded over a day is a noticeable amount. Many physical effects of historical importance, the precession of the equinoxes, the gravitational red shift of light, the slowing of clocks in motion, are of comparable smallness on a per-cycle basis, and become observable only when patiently accumulated. Physicists are people who believe that small effects, properly counted, can move mountains. Engineers are physicists who learned to add a temperature-compensated pendulum.*

HOT Example 22

A simple pendulum of period T = 2.00s swings through an angular amplitude of 2.0° on each side of the vertical. (a) Is the motion strictly simple harmonic? Justify your answer using a numerical comparison. (b) Calculate the maximum speed of the bob. (c) Calculate the height of the bob above its lowest position at the instant of maximum displacement. Take $g = 9.8m/s^2$.

Solution

The angular amplitude of the swing is $\theta_{max} = 2.0° = 2.0\times(\pi/180)rad = 0.034907rad$.

(a) The strict requirement for SHM is $\sin\theta_{max} \approx \theta_{max}$. Compute both:

$$\theta_{max} = 0.034907rad$$
$$\sin\theta_{max} = \sin(2.0°) = 0.034899$$

The difference is 0.000008, a fractional error of about 2×10^{-4}, or 0.02%. The small-angle approximation is therefore correct to two parts in ten thousand at this amplitude, far better than any school stopwatch could detect, and far better than the timing accuracy of the experiment we conducted in Example 19. The motion is strictly simple harmonic to better than the limits of measurement, even though, in mathematical principle, the small-angle approximation introduces a tiny error.

(b) From $T = 2\pi\sqrt{l/g}$, solve for l:

$$l = \frac{gT^2}{4\pi^2} = \frac{(9.8\text{m/s}^2)(2.00\text{s})^2}{4\pi^2}$$

$$l = \frac{9.8 \times 4}{4\pi^2}\text{m} = \frac{9.8}{\pi^2}\text{m} = 0.9928\text{m}$$

The arc-length amplitude (the maximum displacement along the curve of motion) is:

$$A = l\theta_{max} = (0.9928\text{m})(0.034907\text{rad}) = 0.0347\text{m} = 3.47\text{cm}$$

The angular frequency is:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{2.00\text{s}} = \pi\ (\text{rad/s}) = 3.142\text{rad/s}$$

The maximum speed occurs at the lowest point (the equilibrium position) and equals $A\omega$:

$$v_{max} = A\omega = (0.0347\text{m})(3.142\text{rad/s}) = 0.109\text{m/s}$$

So the bob's maximum speed at the bottom of its swing is approximately 0.109m/s, or about 11cm/s.

(c) The height above the lowest position at maximum displacement is:

$$h = l(1 - \cos\theta_{max}) = 0.9928\text{m}(1 - \cos 2.0°) = 6.05 \times 10^{-4}\text{m} \approx 0.6\text{mm}$$

We may **verify the answer** using energy conservation. At maximum displacement all the energy is gravitational potential energy; at the lowest point all the energy is kinetic. Setting them equal:

$$mgh = \frac{1}{2}mv_{max}^2$$

$$h = \frac{v_{max}^2}{2g} = \frac{(0.109\text{m/s})^2}{2 \times 9.8\text{m/s}^2} = 6.07 \times 10^{-4}\text{m}$$

Agreement to three significant figures. The two methods, SHM kinematics and energy conservation, give the same answer, as they must.

Making Sense of the Answer: *A 1-metre pendulum swinging through 2° rises only 0.6mm above its lowest position at the extremes of its arc. This is much smaller than most students expect on first sight; the swing looks much larger to the eye than the height it actually achieves. The reason is geometric: at small angles the bob moves along the arc almost horizontally, with very little vertical component to its displacement. The vertical rise grows as the square of the angular amplitude, and at 2° the square is small indeed. By the time the swing reaches 30°, however, the bob is rising about 13cm, easily visible, and the small-angle approximation is no longer trustworthy.*

Think Like a Physicist: *The two methods agreed to three significant figures, which is the best confirmation we can ask for. SHM kinematics gave us $A = l\theta_{max}$ for the arc amplitude, then ω, then $v_{max} = A\omega$. Energy conservation gave us $h = v_{max}^2/(2g)$ directly from the lowest-point speed. Both routes are valid. When two independent methods give the same number, you know your physics is right. This is one of the secret pleasures of doing physics, the relief of seeing the same answer arrive at the same destination from two different roads.*

We have spent the last several pages on the simple pendulum. The same machine, a restoring force linear in displacement, an SHM equation, a period of the form 2π divided by the square root of a system constant, has appeared once again, but this time the system constant came not from a spring but from gravity acting along an arc. The pattern is now well-established: identify the restoring force, write Newton's second law, identify ω, take $2\pi/\omega$, and the period drops out.

The next three systems we shall study, a column of liquid in a U-tube, a floating cylinder bobbing in water, and a small ball sliding in the bottom of a curved bowl, will each have their own restoring force, their own derivation, and their own ω. None of them involves a spring. None of them involves gravity along an arc in the simple-pendulum sense. And yet, in each case, the same equation $a = -\omega^2 x$ will reappear, and a corresponding period $2\pi/\omega$ will follow. The simple pendulum is the second of eight oscillating systems in this chapter, each different in physical detail and identical in mathematical bone structure.

OTHER STANDARD SHM SYSTEMS: U-TUBE, FLOATING BODY, AND BALL IN A BOWL

We promised at the close of the simple-pendulum section that the next three systems would have nothing in common physically and everything in common mathematically. Now is when we deliver on that promise. A column of liquid in a U-tube has no spring, no string, no gravity acting along an arc; it has only the weight of the unbalanced part of the liquid and a hollow tube to constrain it. A floating cylinder bobbing gently in calm water has no spring either; the restoring force comes from buoyancy alone. A small ball sliding in the bottom of a curved bowl has no string; gravity pulls it back to the bottom because the bowl is curved. And yet, in every one of these three systems, when we work through the geometry slowly and carefully, the same equation $a = -\omega^2 x$ emerges, and the same period $T = 2\pi/\omega$ follows.

Each of the three derivations follows the recipe established in earlier sections: write the picture before the formula, identify the restoring force as a function of small displacement, apply Newton's second law, read off ω^2, and write down the period. The recipe is now well-rehearsed; we shall apply it three times in quick succession.

Liquid oscillations in a U-tube

Imagine a U-shaped glass tube of uniform cross-section **A**, standing vertically and containing a quantity of liquid of density **ρ** that fills both arms to the same equilibrium level. The total length of the liquid column, measured along the centre line of the tube from one free surface, around the bend, to the other free surface, is **L**. When the liquid is at rest, the two free surfaces are at the same height; the pressure at any point in the bend is the same from both sides; and there is no net force pushing the liquid one way or the other.

Now disturb the liquid: blow gently into one arm, or rock the tube to one side, so that the liquid surface in one arm is pushed down by a small distance **y** and the surface in the other arm rises by exactly the same distance **y**. The total mass of liquid is unchanged, of course; the liquid simply rearranges itself between the two arms while maintaining the same total volume. The two free surfaces are now **2y** apart in height.

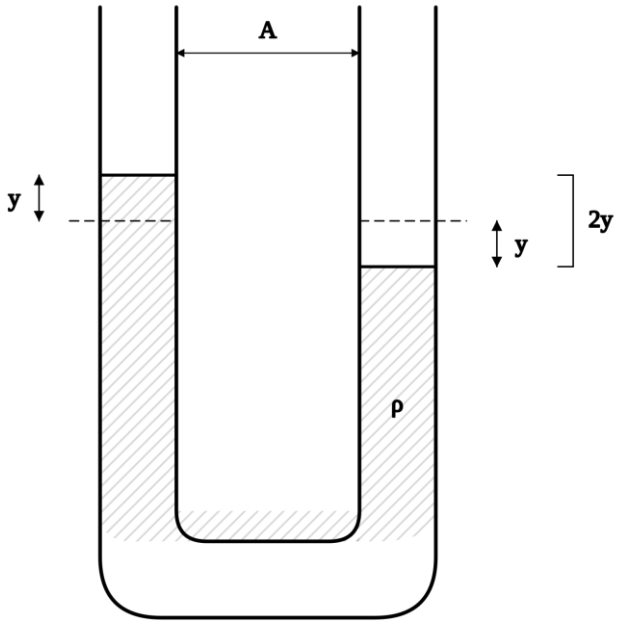


Figure: *A U-tube of uniform cross-section A and density ρ, displaced from equilibrium so that one arm rises by y and the other falls by the same y. The two free surfaces are 2y apart, and the excess weight of the higher column provides the restoring force on the entire liquid mass.*

What is the restoring force? Look only at the column of excess liquid: a column of height 2y in the arm that has been raised, of cross-section A, and density ρ. Its mass is ρA(2y), and its weight pulling downwards on the rest of the liquid is:

$$F = -2y\rho gA$$

The negative sign reflects that this weight always pulls the liquid back towards the equilibrium configuration: when y is positive (one arm raised), the force is downward on the raised side, which means it accelerates the liquid back to equilibrium; when y is negative (that arm lowered, the other raised), the force is in the opposite direction. It is a true restoring force, linear in the displacement, and proportional to y through the constant 2ρgA.

Newton's second law applies to the entire liquid column, of mass $m_{liq} = \rho AL$:

$$\rho AL\frac{d^2y}{dt^2} = -2y\rho gA$$

The mass density and the cross-section both cancel, leaving:

$$\frac{d^2y}{dt^2} = -\frac{2g}{L}y$$

This is the SHM equation, with $\omega^2 = 2g/L$. The period of oscillation is therefore:

$$\mathbf{T = \frac{2\pi}{\omega} = 2\pi\sqrt{\frac{L}{2g}}}$$

It is sometimes more useful to write the period in terms of half the column length, h = L/2 or L=2h, which is the equilibrium height of liquid in each arm above the bend. Substituting:

$$T = 2\pi\sqrt{\frac{h}{g}}$$

Notice the form: it is identical to the period of a simple pendulum, with h in place of l. A liquid column of half-length 1.0m oscillates in a U-tube with the same period as a simple pendulum of length 1.0m. Neither the cross-section of the tube nor the density of the liquid matters; only the length of the liquid column does. Mercury and water in identical tubes oscillate at the same rate.

A floating cylinder bobbing in equilibrium

The next system is a hollow cylindrical tube of mass m and uniform cross-sectional area **A**, weighted at its lower end so that it floats vertically in a liquid of density **ρ** (think of a test tube with a small lead weight at its base, or a cork plumb bob, or a fisherman's float on a calm lake). At equilibrium, the cylinder sits with some depth **h** below the liquid surface; the upthrust on the submerged volume exactly balances the weight of the cylinder.

Setting weight equal to upthrust at equilibrium:

$$mg = \rho Ahg \quad \Rightarrow \quad h = \frac{m}{\rho A}$$

The equilibrium depth h is therefore the mass of the cylinder divided by the product of liquid density and cross-section. It is the depth the cylinder would naturally sink to if released gently from rest just at the surface.

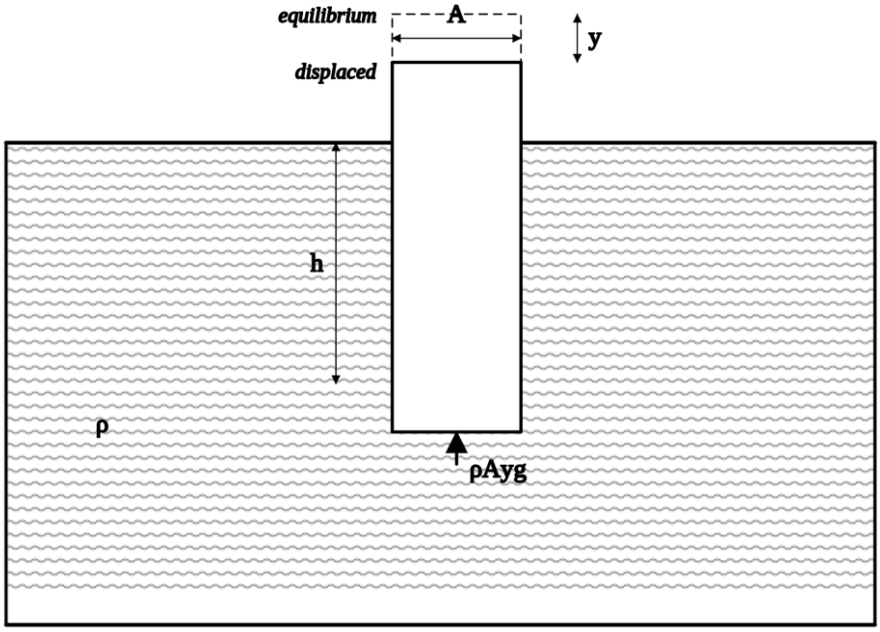


Figure: *A floating cylinder of cross-section A and mass m at its equilibrium submerged depth h, then displaced downward by a further small distance y. The extra submerged volume Ay produces an excess upthrust ρAyg directed upward, which is the restoring force on the cylinder.*

Now push the cylinder downward by a small additional distance y. The submerged volume increases by Ay, and so the upthrust increases by ρAyg, while the weight is unchanged. The net force on the cylinder is now the difference between the new (larger) upthrust and the unchanged weight:

$$F_{net} = \rho A(h+y)g - mg = \rho Ayg$$

Using mg = ρAhg from the equilibrium condition. If we measure y downward as positive, then the force is upward (negative y direction) when y is positive, and downward when y is negative. So:

$$F = -\rho Ayg$$

Newton's second law:

$$m\frac{d^2y}{dt^2} = -\rho Ayg$$

Using m = ρAh from the equilibrium condition, and dividing both sides by m:

$$\frac{d^2y}{dt^2} = -\frac{g}{h}y$$

This is the SHM equation again, with $\omega^2 = g/h$. The period is:

$$T = 2\pi\sqrt{\frac{h}{g}}$$

Once again the same form. $T = 2\pi\sqrt{h/g}$ for the U-tube, $T = 2\pi\sqrt{l/g}$ for the simple pendulum, and now $T = 2\pi\sqrt{h/g}$ for the floating cylinder. The same equation, three times, with three different physical meanings for the length scale h or l. The cross-sectional area, the density of the liquid, and the mass of the cylinder all play their part in setting the equilibrium depth, but once that depth is fixed the period depends only on it and on the local strength of gravity.

A small ball sliding in a curved bowl

The third system is a small smooth ball of radius r resting at the bottom of a polished concave bowl whose inner surface forms part of a sphere of radius R. The bowl is much larger than the ball ($R \gg r$), so when the ball is displaced slightly along the curved surface it slides without rolling, in the way a marble slides over an oiled glass surface. (The case in which the ball rolls without slipping, with friction supplying the rolling constraint, gives a slightly different period and is treated in a later section of this chapter.)

At equilibrium, the ball rests at the lowest point of the bowl, directly below the centre of curvature C. The line from C to the centre of the ball is vertical, of length R - r (the centre of the ball lies on a smaller sphere of radius R - r concentric with the bowl).

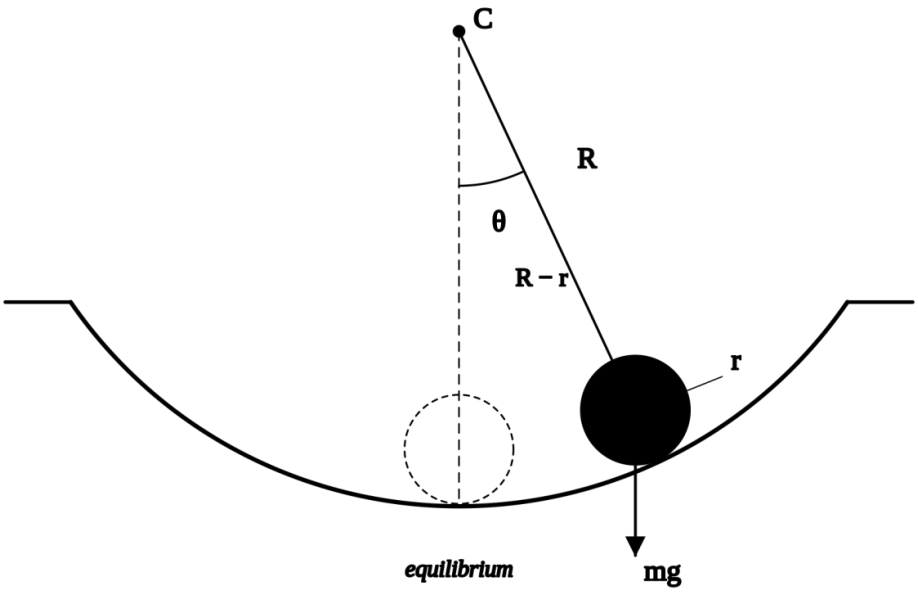


Figure: ***A small ball of radius r sliding in a concave bowl of radius R.*** *The centre of curvature C is at the top, and the ball's centre travels on an arc of radius R − r about C. When displaced through a small angle θ from the vertical at C, the tangential component of the weight along the arc is the restoring force.*

Now displace the ball through a small angle θ from the vertical at C, measured at the centre of curvature. The ball is now resting on the sloping side of the bowl, and the weight mg resolves into two components, exactly as it did for the simple pendulum bob: a component $mg\cos\theta$ along the line from the ball back to C (balanced by the normal force from the bowl), and a component $mg\sin\theta$ tangent to the arc, directed back towards equilibrium. The tangential component is the restoring force.

The arc length of the ball's displacement, measured along the path of the ball's centre, is $s = (R - r)\theta$. For small angles, $\sin\theta \approx \theta$, so the restoring force is:

$$F_t = -mg\sin\theta \approx -mg\theta = -\frac{mg}{R - r}s$$

Apply Newton's second law along the arc:

$$m\frac{d^2s}{dt^2} = -\frac{mg}{R - r}s$$

The mass cancels:

$$\frac{d^2s}{dt^2} = -\frac{g}{R - r}s$$

SHM, with $\omega^2 = g/(R - r)$. The period is:

$$\mathbf{T = 2\pi\sqrt{\frac{R - r}{g}}}$$

Once more the same shape. The role of the length scale is now played by R - r, the distance from the centre of curvature to the centre of the ball. If the ball is much smaller than the bowl, $r \ll R$, and the period reduces to $T \approx 2\pi\sqrt{R/g}$, the same as a simple pendulum of length equal to the bowl's radius.

Four worked examples now follow. Set the kettle on, prepare the calculator, and proceed. The first three are kind to the reader; the fourth is the kind of question that briefly makes a student question their life choices, then redeems itself in the final line.

BINDER Example 23

A wooden cylindrical bar of mass m = 0.40kg and uniform cross-sectional area $A = 12.0\text{cm}^2$ floats vertically in fresh water, slightly weighted at one end so that its long axis stays vertical. The bar is pushed down a small distance and released. (a) Show that its motion is simple harmonic. (b) Calculate the period of oscillation. Take $\rho_{water} = 1000\text{kg/m}^3$ and $g = 9.8\text{m/s}^2$.

Solution

(a) Let the bar have mass m, cross-section A, and float in liquid of density ρ. When displaced downward by a small distance y, the extra submerged volume Ay produces an extra upthrust ρAyg directed upward, while the weight is unchanged. So the net force is restoring:

$$F = -\rho Agy$$

Newton's second law:

$$m\frac{d^2y}{dt^2} = -\rho Agy \quad \Rightarrow \quad a = \frac{d^2y}{dt^2} = -\frac{\rho Ag}{m}y$$

The acceleration is proportional to the displacement (y) and directed back to equilibrium, so the motion is simple harmonic, with angular frequency:

$$\omega = \sqrt{\frac{\rho Ag}{m}}$$

(b) The period follows from $T = 2\pi/\omega$:

$$T = 2\pi\sqrt{\frac{m}{\rho Ag}}$$

Substitute now, with $A = 12.0\text{cm}^2 = 12.0 \times 10^{-4}\text{m}^2$:

$$T = 2\pi\sqrt{\frac{0.40\text{kg}}{(1000\text{kg/m}^3)(12.0 \times 10^{-4}\text{m}^2)(9.8\text{m/s}^2)}} = 1.16\text{s}$$

So the period is approximately 1.16s.

Making Sense of the Answer: *A period of about 1.2s for a 0.40kg wooden bar of cross-section 12cm² in fresh water is comfortable to imagine: roughly one bob per second, of the kind one sees on any calm pond when something has just disturbed a floating reed. Notice that the formula $T = 2\pi\sqrt{m/(\rho Ag)}$ can be rewritten using the equilibrium depth $h = m/(\rho A)$, giving the cleaner form $T = 2\pi\sqrt{h/g}$. This is exactly the form we derived in the prose, and it is identical to the simple pendulum formula with h in place of l.*

Think Like a Physicist: *If the bar were placed in salt water (density 1025kg/m^3 instead of 1000kg/m^3), the period would shorten very slightly, by about 1.2%. If instead the bar were placed in mercury, of density 13600kg/m^3, it would barely sink at all (its equilibrium depth is now only about 2.5cm), and the period would shrink to about 0.31s, roughly three bobs per second. The bar would sit nearly on top of the mercury and twitch very rapidly. The same physics, the same formula, dramatically different numbers.*

BINDER Example 24

A glass U-tube of uniform bore is partly filled with mercury. The total length of the mercury column, measured from one free surface around the bend to the other, is L = 36.0cm. The mercury is given a small displacement and allowed to oscillate freely. (a) Show that the oscillation is simple harmonic. (b) Calculate the period of oscillation. Take $g = 9.8\text{m/s}^2$.

Solution

(a) Let the cross-section be A, the density ρ, and the total liquid length L. When one arm is displaced downward by a small distance y, the other arm rises by the same y, and the excess column of height 2y has weight 2yρgA, which acts as the restoring force on the entire mercury mass:

$$F = -2y\rho gA$$

Newton's second law applied to the whole mercury column, of mass ρAL:

$$\rho AL\frac{d^2y}{dt^2} = -2y\rho gA$$

$$a = \frac{d^2y}{dt^2} = -\frac{2g}{L}y$$

The acceleration is proportional to the displacement (y) and directed back to equilibrium. The motion is simple harmonic, with $\omega^2 = 2g/L$.

(b) The period is:

$$T = 2\pi\sqrt{\frac{L}{2g}}$$

Substitute L = 0.360m and g = 9.8m/s^2:

$$T = 2\pi\sqrt{\frac{0.360\text{m}}{2 \times 9.8\text{m/s}^2}} = 0.852\text{s}$$

So the period of oscillation is approximately 0.85s.

Making Sense of the Answer: *Notice what does and does not appear in the formula $T = 2\pi\sqrt{L/(2g)}$. The cross-section A does not appear; nor does the density ρ. A wide tube and a narrow tube, filled to the same total liquid length, give exactly the same period; mercury and water filled to the same length oscillate at the same rate. The only system property that matters is the total length of the liquid column L. This cancellation of density and cross-section is one of the small surprises of the U-tube.*

Think Like a Physicist: *The form $T = 2\pi\sqrt{L/(2g)}$ can be rewritten as $T = 2\pi\sqrt{h/g}$ with h = L/2, the length of liquid in each arm above the bend. This is the same formula as for the simple pendulum and the floating cylinder; only the geometric meaning of the length scale changes from one system to the next. Whenever you see $T = 2\pi\sqrt{(length)/g}$ written down in any future chapter, it is worth pausing to identify which physical length the formula is referring to.*

REAL Example 25

On a calm dawn on Lake Tanganyika, a fisherman watches his small wooden float bob lazily on the slack of his line. Each time a small fish nibbles, the float dips downward by a centimetre or two and returns. The bobbing is steady, and to the fisherman's ear it has a definite rhythm, perhaps two complete cycles every second.

(a) Explain why does the float bob at this particular rate, and discuss what would make it bob faster or slower?

(b) For a wooden float of mass 30g and cross-sectional area 6cm^2 floating in fresh water, calculate:

(i) the equilibrium submerged depth h,

(ii) the period of oscillation T, and show that the result agrees with the fisherman's observation of about two complete cycles per second.

Solution

The float behaves as a floating cylinder. Its period of small vertical oscillations is:

$$T = 2\pi\sqrt{\frac{h}{g}},$$

where h = m/(ρ A) is the equilibrium submerged depth, the depth the float sinks to when it is left alone in still water under its own weight.

Three things determine h, and therefore T:

1) **The mass of the float, m.** A heavier float pushes more water aside before its weight is balanced, so its equilibrium depth is greater, and its period is longer. A wooden float and a brass cylinder of the same shape would bob at very different rates: the brass cylinder, being denser, would sink to a deeper equilibrium and bob more slowly.
2) **The cross-sectional area of the float at the waterline, A.** A wider float (larger A) needs less depth to displace the same weight of water, so its equilibrium depth is shallower and its period is shorter. A narrow needle-shaped float bobs more slowly than a fat barrel-shaped one of the same mass.
3) **The density of the liquid, ρ.** In denser water, less depth is needed to balance the float's weight, so h is smaller and the period is shorter. The fisherman would notice a tiny difference if he moved his float from the fresh water of Lake Tanganyika to the salt water of the Indian Ocean: the float would bob (about 1%) faster in salt water than in fresh.

(b) (i) The equilibrium depth is:

$$h = \frac{m}{\rho A} = \frac{30 \times 10^{-3}\text{kg}}{(1000\text{kg/m}^3)(6 \times 10^{-4}\text{m}^2)} = 0.050\text{m} = 5.0\text{cm}$$

(ii) And the period is:

$$T = 2\pi\sqrt{\frac{0.050\text{m}}{9.8\text{m/s}^2}} = 2\pi \times 0.0714\text{s} = 0.45\text{s}$$

$$f = \frac{1}{T} = \frac{1}{0.45\text{s}} = 2.22\text{s}^{-1}$$

The frequency is roughly two cycles every second, exactly as the fisherman's ear told him.

Making Sense of the Answer: *The pleasure of this calculation is that it requires no measurement of the spring constant of anything, no estimation of the height the float was pushed down, and no assumption about how hard the fish tugged on the line. Once the fisherman picks his float (fixing its mass and width), he has fixed the period; he has essentially built a slow analogue clock that ticks at* 0.45s *whenever the line is twitched. Different floats give different periods, and an experienced fisherman recognises the 'rhythm' of his particular float so well that he can tell the difference between a fish nibbling and the line catching on a reed by the bobbing rate alone.*

Think Like a Physicist: *The cancellation in $T = 2\pi\sqrt{h/g}$ is the same one we saw in the simple pendulum: the period depends only on a length scale and on the local strength of gravity, not on the inertia or the restoring stiffness separately. The mass and the buoyancy stiffness conspire so that their effects cancel except through the equilibrium depth h. This is the third time we have met this kind of cancellation in this chapter, and it will not be the last. Wherever a system has a natural length scale that is set jointly by inertia and restoring force, the period takes the form $2\pi\sqrt{(length)/g}$ or its non-gravitational analogue.*

HOT Example 26

A hollow cylindrical test tube of uniform cross-sectional area $A = 1.50\text{cm}^2$ and total length $L = 20.0\text{cm}$ is loaded with lead shot at its sealed lower end so that it floats vertically in fresh water with the open end uppermost. The total mass of the tube and its load is M.

(a) Show that the equilibrium submerged depth is $h = M/(\rho A)$, and that the tube bobs in simple harmonic motion of period $T = 2\pi\sqrt{h/g}$.
(b) Find the maximum value M_{max} for which the tube can still float, that is, before it would be fully submerged at equilibrium.
(c) The tube is observed to bob with a period of $T = 0.70\text{s}$. Find M, and verify that $M < M_{max}$.
(d) An additional 5.0g of sand is now poured into the tube. Find the new period of oscillation.
Take $\rho_{water} = 1000\text{kg/m}^3$ and $g = 9.8\text{m/s}^2$.

Solution

(a) At equilibrium, the weight of the loaded tube is balanced by the upthrust on the submerged volume:

$$Mg = \rho Ahg \quad \Rightarrow \quad h = \frac{M}{\rho A}$$

When the tube is displaced downward by a small distance y, the extra submerged volume Ay produces an excess upthrust ρAyg upward, while the weight is unchanged. Newton's second law on the tube (whose total mass remains M) gives:

$$M\frac{d^2y}{dt^2} = -\rho Agy$$

$$a = \frac{d^2y}{dt^2} = -\frac{\rho Ag}{M}y = -\frac{\rho Ag}{\rho Ah}y = -\frac{g}{h}y$$

The motion is therefore simple harmonic, with $\omega^2 = g/h$ and period:

$$T = 2\pi\sqrt{\frac{h}{g}}$$

(b) The tube is fully submerged when h reaches the tube length L. Beyond that, no equilibrium configuration exists with the tube floating; any extra mass and the tube sinks. From $M_{max} = \rho AL$:

$$M_{max} = (1000\text{kg/m}^3)(1.50 \times 10^{-4}\text{m}^2)(0.200\text{m}) = 0.0300\text{kg} = 30.0\text{g}$$

(c) From $T = 2\pi\sqrt{h/g}$, solve symbolically for h then for M:

$$h = \frac{gT^2}{4\pi^2}, \quad M = \rho Ah = \frac{\rho Ag\,T^2}{4\pi^2}$$

Substituting numerical values:

$$M = \frac{(1000\text{kg/m}^3)(1.50 \times 10^{-4}\text{m}^2)(9.8\text{m/s}^2)(0.70\text{s})^2}{4\pi^2} = 0.01825\text{kg} \approx 18.2\text{g}$$

Comparing with the maximum: M = 18.2g is well below M_{max} = 30.0g, so the tube does float, as required.
(d) After adding 5.0g of sand, the new total mass is M' = 18.2g + 5.0g = 23.2g. Rather than recomputing h' from scratch, exploit the proportionality h ∝ M (since h = M/(ρ A) with ρ and A unchanged). The new period is:

$$\frac{T'}{T} = \sqrt{\frac{h'}{h}} = \sqrt{\frac{M'}{M}} = \sqrt{\frac{23.2}{18.2}} = 1.129$$

$$T' = 1.129 \times 0.70\text{s} = 0.79\text{s}$$

So the period lengthens from 0.70s to about 0.79s when the sand is added. (And M' = 23.2g is still below M_{max} = 30.0g, so the tube continues to float.)

Making Sense of the Answer: *The four parts of this question are linked: (a) sets up the formula, (b) finds the upper limit on M, (c) inverts the formula to recover the mass from a measured period, and (d) uses a ratio argument to handle a small change without redoing the full calculation. The ratio trick in (d), $T'/T = \sqrt{M'/M}$, is one of the most useful in oscillator physics; it lets you handle 'what if I change the mass slightly' questions in two lines instead of ten. The numerical answer 0.79s is comfortably close to the original 0.70s, as it should be when the mass changes by less than 30%.*

Think Like a Physicist: *This problem is a textbook example of how a single physical instrument (the loaded test tube) can be used as both a densitometer and a clock. Reading off the period gives you the equilibrium depth and therefore the total mass; reading off the equilibrium depth (with a ruler) gives you the period without ever needing a stopwatch. The instrument is self-calibrating in the sense that h and T are tied to one another by the same constant $4\pi^2/g$, which depends only on the local strength of gravity. Industrial densitometers, used to measure the density of fluids in chemical plants, are loaded test tubes of exactly this kind, instrumented so that the bobbing period can be read electronically to four significant figures. The physics is no more complicated than what you have just done.*

We have analysed five different oscillating systems so far in this chapter. Each system gave us the equation $a = -\omega^2 x$ (or its angular twin), and each gave us a period of the form $2\pi/\omega$ in which the system properties combined to define a single length scale or stiffness scale. All of these systems share a deeper feature, however, that we have so far been able to ignore: the moving body in each case has been treated as a point mass. The ball in the bowl was a small ball; the cylinder bobbing in water was treated as a rigid block whose orientation never changed; the pendulum bob was a point. In every case the body moved without rotating about its own centre of mass.

The next three sections lift this restriction. We shall study three systems in which the oscillating body is a rigid extended object, and in which the rotation of the body about its own axis is an essential part of the motion. The compound pendulum (a real clock pendulum, with a rod and a disc), the torsional pendulum (a body suspended from a wire that twists), and the rolling oscillator (a ball that rolls inside a bowl rather than slides) all require the moment of inertia, the parallel axis theorem, and the radius of gyration that we developed in the *rotation of rigid bodies* chapter of ***Foundations of Physics*** book. If the moment of inertia and the parallel axis theorem are not fresh in your memory, take a moment to revisit the rotation of rigid bodies chapter before reading on.

THE COMPOUND PENDULUM (PHYSICAL PENDULUM)

We promised at the close of the previous section that the next three systems would lift one restriction we had been quietly keeping in the background. In every system so far, the moving body was treated as a point mass. Now we are going to allow the moving body to be a real, extended, rigid object that rotates as it oscillates. The system properties we shall need are:

- the moment of inertia I,
- the parallel axis theorem $I_O = I_{cm} + Md^2$, and
- the radius of gyration k.

These were all developed in the ***Foundations of Physics*** book.

We begin with the most familiar of the three rigid-body oscillators: the **compound pendulum**, sometimes called the **physical pendulum**. The compound pendulum is the workhorse of every clockmaker's bench, every metronome, every grandfather clock face, and every laboratory experiment in which a metre rule with holes drilled along its length is hung from a pin and timed. It is not a simple pendulum, although its mathematics is built on top of the simple pendulum's, and its formula reduces to the simple pendulum's formula in a special limit we shall identify shortly.

What a compound pendulum is

A compound pendulum, sometimes called a physical pendulum, *is any rigid body that is free to swing about a fixed horizontal axis that does not pass through its centre of mass.* The axis is called the pivot, and the line from the pivot to the centre of mass is the body's effective swing arm. When the body is displaced through a small angle from its equilibrium position (in which the centre of mass hangs directly below the pivot), gravity supplies a restoring torque that pulls the body back, and the body oscillates.

The simple pendulum is the special case in which the swinging body is a point mass at the end of a light string. The string itself has no mass and contributes no moment of inertia; the entire mass of the system is concentrated at a single point at distance l from the pivot. The compound pendulum drops both of these idealisations: *the body has finite size, its mass is distributed throughout its volume, and the body's rotation about its own centre of mass during the swing must be accounted for.*

Examples of compound pendulums include: a metre ruler hung from a hole drilled near one end; the disc-and-rod assembly of a clock pendulum; a uniform bar pivoted at one end and allowed to swing; a wheel hung from a peg on its rim; even an athlete on a high bar at the moment they are fully extended below the bar before swinging up. All of these are real bodies of finite extent rotating about real axes that do not pass through their centres of mass, and all of them obey the same equation we are about to derive.

From restoring torque to the period formula

Consider a rigid body of mass M suspended from a fixed horizontal pivot at point O. The centre of mass of the body is at point C, at distance d from the pivot. At equilibrium, the line from O to C is vertical, and C hangs directly below the pivot. Now displace the body through a small angle θ from the vertical, as shown in the figure.

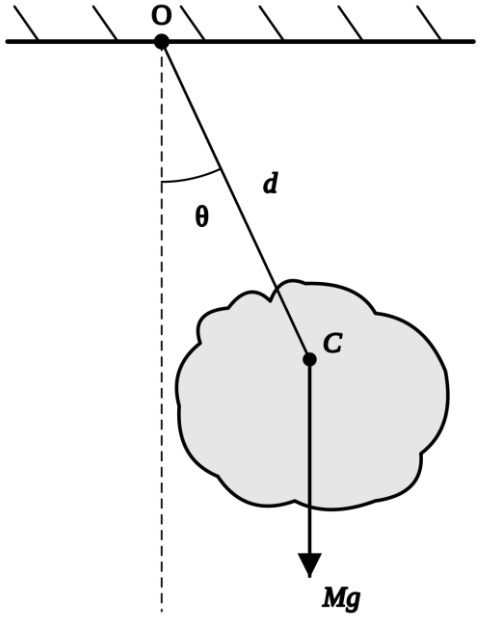


Figure: ***A rigid body of mass M suspended from a fixed pivot O.*** *The centre of mass C is at distance d from O. Displaced by a small angle θ from the vertical, the body experiences a restoring torque produced by the weight Mg acting at C.*

The only external forces on the body are its weight Mg (acting at C) and the reaction at the pivot (which exerts no torque about O, since its lever arm is zero). The torque about the pivot is therefore due entirely to gravity:

$$\tau = -Mg \times d\sin\theta$$

The negative sign captures the restoring nature of the torque: when θ is positive (body tilted to the right of vertical), the gravity torque rotates the body anticlockwise about O (back toward equilibrium), which is a negative torque in our sign convention. The factor d is the lever arm: the perpendicular distance from the pivot to the line of action of the weight (which is vertical at C, so the lever arm is dsin θ).

Newton's second law for rotation about the pivot states:

$$\tau = I_O\alpha = I_O\frac{d^2\theta}{dt^2}$$

where I_O is the moment of inertia of the body about the pivot (not about the centre of mass), and $\alpha = d^2\theta/dt^2$ is the angular acceleration. Equating the two expressions for the torque:

$$I_O\frac{d^2\theta}{dt^2} = -Mgd\sin\theta$$

For small angles, $\sin\theta \approx \theta$, and the equation reduces to:

$$\alpha = \frac{d^2\theta}{dt^2} = -\frac{Mgd}{I_O}\theta$$

This has the standard SHM form $d^2\theta/dt^2 = -\omega^2\theta$, with:

$$\omega^2 = \frac{Mgd}{I_O}$$

And the period of small oscillations is therefore:

$$\mathbf{T = \frac{2\pi}{\omega} = 2\pi\sqrt{\frac{I_O}{Mgd}}}$$

This is the master formula for the compound pendulum. Three quantities determine the period: the moment of inertia about the pivot, I_O; the mass of the body, M; and the distance from the pivot to the centre of mass, d (along with the local strength of gravity g).

If the body were a point mass at distance d from the pivot, then $I_O = Md^2$, and the formula reduces to $T = 2\pi\sqrt{d/g}$, the simple pendulum formula. The compound pendulum is therefore a generalisation of the simple pendulum, and the simple pendulum is a special case of the compound.

The parallel-axis form, and the radius of gyration

It is often more convenient to express I_O in terms of the moment of inertia about the centre of mass, I_{cm}, which is usually a tabulated property of the body. The parallel axis theorem tells us how:

$$I_O = I_{cm} + Md^2$$

Substituting into the period formula:

$$T = 2\pi\sqrt{\frac{I_{cm} + Md^2}{Mgd}}$$

If we define the radius of gyration of the body about its centre of mass by $I_{cm} = Mk^2$. Then k has the dimensions of length, and the formula simplifies further:

$$\mathbf{T = 2\pi\sqrt{\frac{Mk^2 + Md^2}{Mgd}} = 2\pi\sqrt{\frac{k^2 + d^2}{gd}}}$$

Notice that the mass M has dropped out completely. The period of a compound pendulum depends on the geometry of the mass distribution (through k and d) and on gravity, but not on the total mass at all. This is the same isolation of mass that we saw for the simple pendulum, and for exactly the same reason: *heavy bodies feel a larger gravitational torque, but they also have larger inertia to be turned, and the two effects cancel.* Two compound pendulums of identical shape but different masses (one made of wood, the other of brass, with the same dimensions) oscillate at exactly the same period.

An equivalent form of this formula, useful for experimental work, is:

$$\mathbf{T^2 d = \frac{4\pi^2(d^2 + k^2)}{g}}$$

This form is linear in d^2 when $T^2 d$ is plotted against d^2, with slope $4\pi^2/g$ and intercept $4\pi^2 k^2/g$. By measuring T at several pivot positions d on the same body and plotting the graph, one extracts both g (from the slope) and k (from the intercept) from a single run of measurements. The perforated metre rule experiment in the laboratory works on exactly this principle.

The equivalent simple pendulum

Compare the compound pendulum period $T = 2\pi\sqrt{(k^2+d^2)/(gd)}$ with the simple pendulum period $T = 2\pi\sqrt{l/g}$. For the two to give the same period, we must have $l = (k^2+d^2)/d$. This length L_{eq} is called the **equivalent simple pendulum length**:

$$\mathbf{L_{eq} = \frac{k^2 + d^2}{d} = d + \frac{k^2}{d}}$$

It is *the length of a simple pendulum that would oscillate at exactly the same period as the given compound pendulum.* Notice the structure: L_{eq} is the sum of two terms, d (the distance from pivot to centre of mass) and k^2/d (a correction term that captures the body's extended mass distribution). Since both terms are positive, $L_{eq} > d$ always; the equivalent simple pendulum is always longer than the lever arm of the gravity torque. A compound pendulum therefore swings more slowly than a simple pendulum of the same lever arm d; the extra mass distributed throughout the body acts as additional rotational inertia and slows the swing. Only in the limit $k \to 0$ (a body whose mass is concentrated at a single point at distance d from the pivot, a simple pendulum) does L_{eq} reduce to d.

The minimum period

The expression $L_{eq} = d + k^2/d$ has an interesting feature: it depends on d, and not in a monotonic way. If the pivot is moved very close to the centre of mass ($d \to 0$), the term $k^2/d \to \infty$, and the period becomes very long; the body barely swings, since

the gravity torque has a vanishing lever arm. If the pivot is moved very far from the centre of mass ($d \to \infty$), the term d itself dominates and L_{eq} grows without bound; the body swings slowly because the equivalent length is now very large. Somewhere between these two extremes there must be a value of d that minimises the period. To find it, differentiate L_{eq} with respect to d and set the result to zero:

$$\frac{dL_{eq}}{dd} = 1 - \frac{k^2}{d^2} = 0 \quad \Rightarrow \quad \mathbf{d = k}$$

The minimum period therefore occurs when the pivot is placed at a distance equal to the radius of gyration from the centre of mass. At this special pivot position:

$$L_{eq,\,min} = k + \frac{k^2}{k} = 2k$$

$$\mathbf{T_{min} = 2\pi\sqrt{\frac{2k}{g}}}$$

This is a universal result: it applies to any rigid body, of any shape. Watchmakers and pendulum-clock designers historically used this result to set the optimum balance between size and speed: a clock pendulum is positioned to give a period close to the minimum for its mass distribution, so that small adjustments to the pivot position produce only second-order changes in the period.

Five worked examples now follow. Put the kettle on, pick up your pen, and proceed. The fifth one is the kind of question that briefly makes a student wish they had not removed the disc from the wall clock at the beginning of the section, then rewards the persistence with a small physical insight.

BINDER Example 27

A uniform rod of mass M = 1.20kg and length L = 1.50m is suspended from one end and allowed to swing as a compound pendulum about that end. (a) Calculate the moment of inertia of the rod about the pivot. (b) Calculate the period of small oscillations. (c) Calculate the equivalent simple pendulum length and explain why it is not equal to the rod length, nor to the distance from the pivot to the centre of mass. Take $g = 9.8\text{m/s}^2$.

Solution

For a uniform rod pivoted at one end, the moment of inertia about the pivot is the standard tabulated result $I_O = ML^2/3$. The centre of mass of a uniform rod is at its midpoint, so d = L/2.

(a) Substitute M = 1.20kg and L = 1.50m:

$$I_O = \frac{ML^2}{3} = \frac{(1.20\text{kg})(1.50\text{m})^2}{3} = \frac{2.70\text{kgm}^2}{3} = 0.900\text{kgm}^2$$

(b) Substitute into $T = 2\pi\sqrt{I_O/(Mgd)}$:

$$T = 2\pi\sqrt{\frac{0.900\text{kgm}^2}{(1.20\text{kg})(9.8\text{m/s}^2)(0.750\text{m})}} = 2.01\text{s}$$

Or, using the symbolic shortcut $T = 2\pi\sqrt{2L/(3g)}$ (which one obtains by substituting $I_O = ML^2/3$ and d = L/2 directly into the master formula):

$$T = 2\pi\sqrt{\frac{2 \times 1.50\text{m}}{3 \times 9.8\text{m/s}^2}} = 2.01\text{s}$$

(c) The equivalent simple pendulum length is:

$$L_{eq} = \frac{I_O}{Md} = \frac{0.900\text{kg} \cdot \text{m}^2}{(1.20\text{kg})(0.750\text{m})} = 1.00\text{m}$$

Or symbolically: $L_{eq} = (ML^2/3)/(M \cdot L/2) = 2L/3 = (2/3)(1.50\text{m}) = 1.00\text{m}$.

This is two-thirds of the rod length, longer than the lever arm d = 0.75m (the distance from the pivot to the centre of mass) but shorter than the full rod length L = 1.50m.

The reason L_{eq} is not equal to d is that the rod is an extended body whose mass is distributed along its length, not concentrated at the centre of mass; the rod has its own moment of inertia about its centre of mass ($I_{cm} = ML^2/12$), which adds to the gravitational lever-arm contribution and slows the swing.

The reason L_{eq} is not equal to L is that the bottom end of the rod is moving faster (greater radius from pivot) than the top end during the swing, so the body's effective inertia is less than what one would calculate by treating the entire rod mass as concentrated at the bottom end.

Making Sense of the Answer: *A 1.50m rod swings with the same period as a simple pendulum of length 1.00m, which has period $T = 2\pi\sqrt{1.00/9.8} \approx 2.01s$, exactly what we found above. The result $L_{eq} = 2L/3$ for a rod pivoted at one end is one of the most useful in compound-pendulum physics; it converts a question about a swinging bar into a question about a simple pendulum, which is much easier to think about. It is worth memorising.*

Think Like a Physicist: *The factor of 2/3 comes from a small piece of integration: the moment of inertia of a rod about its end is $\int_0^L (m/L)x^2dx = ML^2/3$, and the centre of mass is at L/2. The ratio $I_O/(Md) = 2L/3$ falls out of the integration alone. Whenever you see a simple, clean coefficient like 2/3 in a compound-pendulum formula, suspect it of having come from an integration of x^2 over a uniform mass distribution. The whole machinery of moment-of-inertia integrals is always quietly behind the curtain.*

BINDER Example 28

A uniform metre rule with holes drilled at intervals along its length is suspended on a horizontal pin so that it can swing as a compound pendulum. When the pin is placed in the hole at distance d = 30.0cm from the centre of the rule, the period of small oscillations is observed to be T = 1.53s. The radius of gyration of a uniform metre rule about its centre is k = 0.289m. Use these data to determine the local value of g at the laboratory.

Solution

The compound-pendulum period in the radius-of-gyration form is $T = 2\pi\sqrt{(k^2+d^2)/(gd)}$, which rearranges to:

$$g = \frac{4\pi^2(k^2+d^2)}{T^2 d}$$

Substituting:

$$g = \frac{4\pi^2\left((0.289\text{m})^2 + (0.300\text{m})^2\right)}{(1.53\text{s})^2(0.300\text{m})} = 9.76\text{m/s}^2$$

So the local value of g is 9.76m/s^2, which is consistent with the accepted value of 9.8m/s^2 at sea level.

Making Sense of the Answer: *The compound-pendulum experiment is one of the most accurate gravimetric methods accessible in a school laboratory. Unlike the simple pendulum (which requires careful measurement of the string length and assumes a point mass at the end), the compound pendulum uses a single rigid body whose dimensions are easy to measure and whose moment of inertia is built in. A single experimental run at one pivot position, like the one above, gives* g *to within about a percent. A more refined experiment varies the pivot position over several values of* d *and plots* T^2d *against* d^2*; the resulting straight-line slope* $4\pi^2/g$ *can be measured to four significant figures, giving* g *correspondingly accurately.*

Think Like a Physicist: *The radius of gyration k = 0.289m for a metre rule is itself a clean derived result: for a uniform rod of length L about its midpoint, $I_{cm} = ML^2/12$, so $k = \sqrt{I_{cm}/M} = L/(2\sqrt{3})$. For L = 1.000m, this gives $k = 1/(2\sqrt{3})m = 0.2887m$, which we have rounded to 0.289m in the calculation. The same formula structure $k = L/(2\sqrt{3})$ applies to every uniform rod regardless of mass; only the length matters. This independence of mass is one of the recurring small pleasures of moment-of-inertia calculations.*

BINDER Example 29

A uniform circular disc of mass M = 1.50kg and radius R = 0.200m is suspended from a horizontal axis through a point on its rim, perpendicular to the plane of the disc. Calculate the period of small oscillations. Take $g = 9.8\text{m/s}^2$.

Solution

For a uniform disc, the moment of inertia about its centre of mass perpendicular to its plane is the standard result $I_{cm} = MR^2/2$. The pivot is at the rim, so d = R (distance from rim to centre is one radius). By the parallel axis theorem:

$$I_O = I_{cm} + Md^2 = \frac{MR^2}{2} + MR^2 = \frac{3MR^2}{2}$$

Substitute into $T = 2\pi\sqrt{I_O/(Mgd)}$:

$$T = 2\pi\sqrt{\frac{\frac{3MR^2}{2}}{MgR}} = 2\pi\sqrt{\frac{3R}{2g}} = 2\pi\sqrt{\frac{3\times0.200\text{m}}{2\times9.8\text{m/s}^2}} = 1.10\text{s}$$

So the period of small oscillations is approximately 1.10s.

Making Sense of the Answer: *A disc of 1.5kg and 20cm radius swung from its rim oscillates with about a one-second period. The factor* 3R/(2g) *inside the square root is general for a disc-at-rim, regardless of the mass: the period depends only on the radius of the disc and the local strength of gravity. The mass cancellation is by now a familiar feature of compound-pendulum problems.*

Think Like a Physicist: *Notice the role of the parallel axis theorem in this calculation. Without it, computing* I_O *by direct integration over the disc with the rim as the axis would have been a multi-step exercise in cylindrical coordinates. With it, the work collapses to one line:* $I_O = I_{cm} + Md^2$. *This is the recurring strategy in compound-pendulum physics: write down* I_{cm} *from a table, add* Md^2, *and you have* I_O *for any pivot position you like. The same body, the same body, with seven different pivot positions, gives seven different periods, and the parallel axis theorem produces all seven from a single* I_{cm}.

REAL Example 30

Kipanga and Kipute are watching the wall clock in Mr. Akilikubwa's laboratory. The pendulum, a thin brass rod ending in a heavy circular disc, swings back and forth with a steady rhythm. Kipanga turns to Kipute with a furrowed brow.

Kipanga: *"That clock pendulum is basically a heavy disc at the end of a thin rod, so it is almost a simple pendulum. Why does it not swing at the same rate as a simple pendulum of the same length?"*

Kipute thinks for a moment.

Kipute: *"Because a simple pendulum has all its mass at one point. The clock pendulum has mass spread out. The disc at the bottom has its own moment of inertia about the pivot, on top of the lever-arm moment of inertia* Md^2, *and the rod has mass too. The total moment of inertia about the pivot is larger than* Md^2 *alone, so the period is longer."*

Kipanga: *"So the disc makes it slower?"*

Kipute: *"Not by much, because the disc's own moment of inertia about its centre of mass is* $MR^2/2$, *with the disc's radius R typically much smaller than its distance d from the pivot. The total moment of inertia is* $Md^2 + MR^2/2$, *and the second term is a small correction to the first. The period is correspondingly close to the simple-pendulum period of length d, but slightly larger."*

Kipanga: *"But Mr. Akilikubwa said the clock is designed so the period is exactly two seconds, so the seconds-hand ticks once per swing."*

Kipute: *"Yes! That is the whole point. The clockmaker adjusts the position of the disc on the rod to fine-tune the period. Moving the disc up shortens d and slightly changes the moment of inertia at the same time. There is a small adjustment screw at the bottom of every clock pendulum that does exactly this. The clockmaker tunes it once, and the clock keeps time forever, until the temperature changes or someone bumps it."*

Mr. Akilikubwa, who has been listening, walks over and adds:

Mr. Akilikubwa: *"And the beauty is that for small oscillations, the period is independent of amplitude. So whether the clock is wound up enough to swing through three degrees or one degree, the period is the same. That is the entire reason the pendulum clock was the most accurate timekeeping device in the world for nearly three centuries, from 1656 until quartz crystals took over in the 1930s. Galileo noticed the property; Huygens built the first working clock; Harrison perfected it for marine navigation; and the pendulum clock survived for three hundred years before electronics replaced it."*

Kipanga: *"What if the clock pendulum swings hard, with a large amplitude?"*

Mr. Akilikubwa: *"Then the small-angle approximation fails,* $\sin\theta$ *is no longer equal to* θ, *and the period does begin to depend on amplitude. The clock would run slow. That is why clock pendulums are designed to swing through very small arcs, typically less than two degrees on each side of vertical. A wall clock you can see from across the room has a noticeable swing, but if you measure it, it is almost always between one and three degrees. The visible motion is much smaller than it appears."*

Making Sense of the Answer: *The dialogue captures three ideas that are easy to miss on first reading.* ***First:*** *the compound pendulum is slower than the simple pendulum of the same lever arm, because the body's own moment of inertia adds to the lever-arm moment of inertia.* ***Second:*** *clockmakers tune the period by adjusting the disc position, which changes both d and the moment of inertia at the same time.* ***Third:*** *the isochronous property of small oscillations is the reason pendulum clocks worked at all; a clock whose period depended on amplitude would lose time whenever the amplitude changed.*

Think Like a Physicist: *Pendulum clocks ruled timekeeping for three centuries, from Christiaan Huygens's first design in 1656 to the rise of quartz crystals in the 1930s. The key innovation was the escapement mechanism, which delivers a small impulse to the pendulum at each swing to compensate for air resistance and friction at the pivot, keeping the amplitude (and therefore the period) approximately constant over time. The compound pendulum at the heart of every grandfather clock is exactly the system we have just analysed; its period is set by the geometry of the rod and disc, and its accuracy is set by the quality of the escapement. When you next walk past a working pendulum clock, in a sitting room, in a railway station, in a museum, pause to admire it. It is three hundred years of compound-pendulum physics, distilled into a piece of furniture.*

HOT Example 31

Consider again the uniform disc of Example 29: M = 1.50kg, R = 0.200m. The disc is now to be hung from a series of pivot positions along a radial line from the centre to the rim. (a) Find the pivot distance d (measured from the centre of the disc) that gives the minimum period of small oscillations. (b) Calculate this minimum period T_{min}, and compare it with the period

at the rim (Example 29). (c) Explain physically why the period at the rim is close to (but slightly larger than) the minimum, even though the rim is not the optimum pivot. Take $g = 9.8\text{m/s}^2$.

Solution

(a) The minimum period occurs at $d = k$.

The radius of gyration about the centre of mass is defined by $I_{cm} = Mk^2$. For a uniform disc, $I_{cm} = MR^2/2$, so:

$$k = \sqrt{\frac{I_{cm}}{M}} = \sqrt{\frac{R^2}{2}} = \frac{R}{\sqrt{2}}$$

Substitute $R = 0.200\text{m}$:

$$k = \frac{0.200\text{m}}{\sqrt{2}} = 0.1414\text{m} \approx 0.141\text{m}$$

So the optimum pivot is at distance $d_{opt} = k = 0.141\text{m}$ from the centre of the disc.

The pivot distance d that gives the minimum period is approximately 0.141m.

(b) The minimum period is $T_{min} = 2\pi\sqrt{2k/g}$:

$$T_{min} = 2\pi\sqrt{\frac{2 \times 0.1414\text{m}}{9.8\text{m/s}^2}} = 1.067\text{s}$$

Compare with the rim-pivot period from Example 29: $T_{rim} = 1.10\text{s}$. The ratio is:

$$\frac{T_{rim}}{T_{min}} = \frac{1.10}{1.067} = 1.031$$

The rim-pivot period is therefore only about 3% longer than the minimum, even though the rim is not the optimum pivot.

(c) The function $L_{eq}(d) = d + k^2/d$ has a minimum at $d = k$ and is a relatively flat function of d in the neighbourhood of this minimum (because the derivative $1 - k^2/d^2$ is small near $d = k$). For the disc, the rim is at $d = R = 0.200\text{m}$ while the optimum is at $d = k = 0.141\text{m}$. These two pivot positions are close to each other (the ratio $d/k = R\sqrt{2}/R = \sqrt{2} \approx 1.41$, i.e., the rim is only about 40% farther from the centre than the optimum). The flatness of $L_{eq}(d)$ near its minimum then ensures that the period at the rim is only a few percent away from the minimum, even though the pivot positions differ by 40% in absolute terms.

Making Sense of the Answer: *The condition **d = k** for minimum period is **universal**: it applies to any rigid body, of any shape, as long as the body has a definite radius of gyration about its centre of mass. The same condition applies whether the body is a rod, a disc, a hollow sphere, an irregular blob; only the value of k changes.*

Think Like a Physicist: *The flatness of the period function near its minimum is what makes the compound pendulum a good gravimetric instrument. If the period varied sharply with pivot position, small experimental errors in setting up the pivot would translate into large errors in the period and therefore in the inferred g. Because the period varies slowly with d near the minimum, the experimentalist who happens to set the pivot a few millimetres off the optimum still measures a period within a small fraction of the true minimum, and can therefore extract g accurately.*

We have spent the last several pages on the compound pendulum. The same recipe (find the restoring agent, write Newton's second law, identify ω^2, take $2\pi/\omega$) has produced once again a clean SHM equation, this time for a rigid body rotating about a fixed axis. The restoring agent was gravity, acting on the centre of mass through the lever arm d. The inertia was I_O, computed via the parallel axis theorem from the body's own I_{cm} and the pivot displacement.

The next system, the torsional pendulum, replaces gravity as the restoring agent with a different mechanism entirely: the elastic resistance of a wire that has been twisted out of its natural shape.

TORSIONAL OSCILLATION OF A SOLID SPHERE AND CYLINDER

The compound pendulum used gravity as its restoring agent. The torsional pendulum, replaces gravity with a different mechanism altogether: the elastic resistance of a wire that has been twisted out of its natural shape. Hang a body from a vertical wire, twist the body about the wire's axis through a small angle, and the wire fights back with a restoring torque proportional to the twist. The mathematical structure will be familiar by now: identify the restoring agent, write Newton's second law for rotation, identify ω^2, take $2\pi/\omega$, and the period drops out.

The torsional pendulum is, in a precise sense, the rotational twin of the mass-on-a-spring system. Where the mass-spring system has a **stiffness k** (*force per unit displacement*), the torsional pendulum has a **torsional constant c** (*torque per unit angular displacement*). Where the mass-spring system has **inertia m** (*mass*), the torsional pendulum has **inertia I** (*moment of inertia about the wire axis*). ***Wherever the mass-spring formula has m or k, the torsional pendulum formula has I or c.*** The translation is mechanical and the resulting period formula is the **rotational twin** of $T = 2\pi\sqrt{m/k}$.

The torsional pendulum

A torsional pendulum, in its simplest form, consists of a rigid body suspended from a fixed point by a thin vertical wire. The wire passes through the body's centre of mass (or close to it) and along the body's axis of symmetry. When the body hangs at rest, the wire is straight and untwisted; this is the equilibrium configuration. Now grasp the body and rotate it through a small angle θ about the wire's axis. The wire, attached to the body at the bottom and to a fixed support at the top, is now twisted along its length. Release the body, and the elastic resistance of the twisted wire rotates it back. The body overshoots the equilibrium, twists the wire the other way, and oscillates. This is the torsional pendulum, and we shall now derive its period.

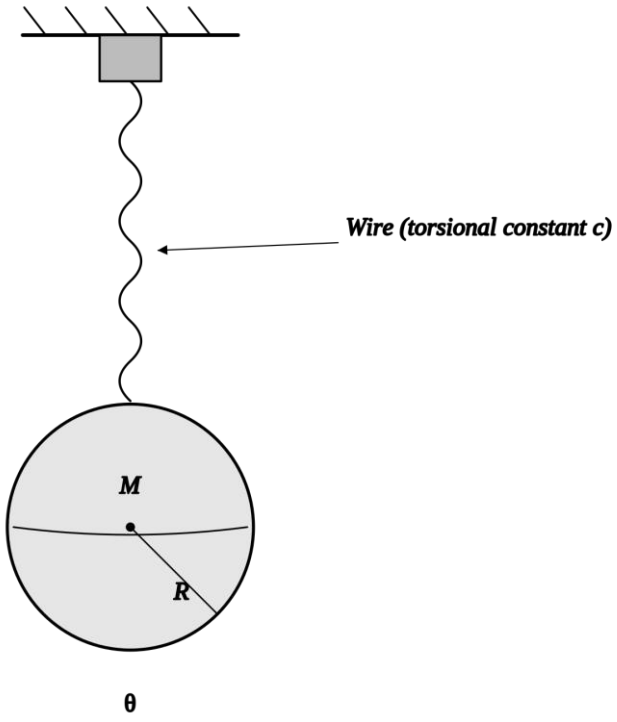


Figure: *A solid body of mass M and moment of inertia I about the wire axis suspended from a vertical wire of torsional constant c. The wire is fixed at the upper end to a small support attached to the ceiling. When the body is twisted through a small angle θ about the wire axis, the wire provides a restoring torque proportional to θ, and the body oscillates.*

restoring torque to the period formula

When the wire is twisted through a small angle θ, it exerts on the body a restoring torque proportional to that angle. This is Hooke's law for torsion, and it holds as long as the wire remains within its elastic limit. Calling the proportionality constant c (the torsional constant of the wire, with units of newton-metres per radian):

$$\tau = -c\theta$$

The negative sign captures the restoring nature of the torque: when θ is positive, τ is negative, and the wire rotates the body back toward equilibrium.

Applying Newton's second law for rotation about the wire axis:

$$I\frac{d^2\theta}{dt^2} = -c\theta$$

Or, dividing both sides by I:

$$\alpha = \frac{d^2\theta}{dt^2} = -\frac{c}{I}\theta$$

This is the SHM equation in angular form, with:

$$\omega^2 = \frac{c}{I}, \qquad \mathbf{T = 2\pi\sqrt{\frac{I}{c}}}$$

Compare with the mass-spring formula $T = 2\pi\sqrt{m/k}$. The replacements are m by I and k by c. Every algebraic feature of the mass-spring system has its rotational twin in the torsional pendulum: the restoring force becomes a restoring torque, the

displacement becomes an angular displacement, the spring constant becomes a torsional constant, the mass becomes a moment of inertia. The mathematical structure is identical, and the physical phenomena are correspondingly identical: oscillations at a single natural frequency, isochronous motion, conservation of energy in the undamped case, sensitivity to the system's inertial and elastic properties. All the intuition you developed for the mass-spring system applies here, with the dictionary translated.

Period for a solid sphere and a solid cylinder

The two most common bodies one suspends from a torsion wire in the laboratory are a uniform solid sphere and a uniform solid cylinder, each rotated about an axis through its centre that coincides with the wire. Their moments of inertia about this axis are standard tabulated results:

$$I_{sphere} = \frac{2}{5}MR^2, \qquad I_{cylinder} = \frac{1}{2}MR^2$$

Substituting into $T = 2\pi\sqrt{I/c}$:

$$T_{sphere} = 2\pi\sqrt{\frac{2MR^2}{5c}}, \qquad T_{cylinder} = 2\pi\sqrt{\frac{MR^2}{2c}}$$

The ratio of the two periods, for the same wire and equal mass and radius, is:

$$\frac{T_{cylinder}}{T_{sphere}} = \sqrt{\frac{I_{cylinder}}{I_{sphere}}} = \sqrt{\frac{1/2}{2/5}} = \frac{\sqrt{5}}{2} \approx 1.118$$

So the cylinder oscillates about 12% more slowly than the sphere of the same mass and radius. The reason is that for the same total mass, the cylinder has more of its mass at large radii from the axis (since a disc-cross-section is uniform-radius all the way out, while a sphere-cross-section is widest at the equator and narrows toward the poles). More mass at large radius means larger moment of inertia, means longer period.

Measuring the torsional constant

The torsional constant c of a wire is rarely tabulated; it depends on the wire's material, its length, its cross-sectional radius, and its surface treatment. It is therefore standard practice in the laboratory to measure c of a given wire by suspending a body of known moment of inertia from it, timing the oscillations, and rearranging $T = 2\pi\sqrt{I/c}$:

$$\mathbf{c = \frac{4\pi^2 I}{T^2}}$$

A solid sphere of accurately known mass and radius is the body of choice, because its moment of inertia $(2/5)MR^2$ is one of the cleanest of all moments of inertia, and any small geometric imperfection in the sphere has only a quadratic effect on I (and therefore only a quadratic effect on the inferred c). Once c is known for a particular wire, the same wire can be used to measure the moment of inertia of any other body, just by hanging the body from the wire, timing the oscillations, and inverting:

$$\mathbf{I = \frac{cT^2}{4\pi^2}}$$

The torsional pendulum is, in this sense, a rotational analogue of the spring balance: just as a spring of known stiffness lets you measure mass by reading the static extension, a wire of known torsional constant lets you measure moment of inertia by reading the period of small oscillations. The first instrument is gravimetric, the second is rotational; the underlying physics is the same.

Two worked examples now follow.

BINDER Example 32

A solid brass sphere of mass M = 2.50kg and radius R = 0.0500m is suspended from a vertical wire fixed to the ceiling. When the sphere is twisted through a small angle and released, it is observed to complete 8 full torsional oscillations in 20.0s. (a) Find the period of oscillation. (b) Calculate the moment of inertia of the sphere about the wire axis. (c) Determine the torsional constant of the wire.

Solution

(a) The period is the total time divided by the number of oscillations:

$$T = \frac{t}{n}$$

$$T = \frac{20.0s}{8} = 2.50s$$

(b) For a uniform solid sphere, the moment of inertia about an axis through its centre is:

$$I = \frac{2}{5}MR^2$$

$$I = \frac{2}{5}(2.50\text{kg})(0.0500\text{m})^2 = 2.50 \times 10^{-3}\text{kgm}^2$$

(c) The torsional constant follows from $T = 2\pi\sqrt{I/c}$:

$$c = \frac{4\pi^2 I}{T^2}$$

$$c = \frac{4\pi^2(2.50 \times 10^{-3}\text{kgm}^2)}{(2.50\text{s})^2} = 1.58 \times 10^{-2}\text{Nm/rad}$$

Making Sense of the Answer: *A torsional constant of ~ 16mNm/rad is what one expects from a thin steel wire about a metre long; it tells us that the wire applies a torque of about 16 millinewton-metres for each radian of twist. If we doubled the twist angle, we would feel double that resistance, in the wire's elastic regime. The brass sphere, with its small moment of inertia of 2.5 × 10⁻³ kgm², produces a period of 2.5s, which is comfortable for a stopwatch and gives clean timing data without requiring particularly fine instruments.*

Think Like a Physicist: *The clean numerical separation between the moment of inertia (~ 10^{-3}kgm^2) and the torsional constant (~ 10^{-2}Nm/rad) is a feature, not a bug. The two quantities differ by a factor of ten in their natural orders of magnitude, which is what makes the period (which depends on the ratio* I/c*) come out to a few seconds rather than a few microseconds or a few hours. In molecular vibrations or MEMS (Micro-Electro-Mechanical Systems) oscillators, the same formula gives periods of nanoseconds, simply because* I *in those systems is many orders of magnitude smaller than in our brass sphere.*

BINDER Example 33

The brass sphere of Example 32 is now removed from the wire and replaced by a solid brass cylinder of the same mass M = 2.50kg and the same radius R = 0.0500m, suspended so that the cylinder's axis of symmetry coincides with the wire. Without retiming, predict the new period and verify that it is about 12% longer than the period of the sphere.

Solution

The wire is unchanged, so its torsional constant c is the same. The mass and radius are unchanged. The only thing that changes is the moment of inertia:

$$I_{\text{sphere}} = \frac{2}{5}MR^2, \qquad I_{\text{cylinder}} = \frac{1}{2}MR^2$$

The ratio of the new to the old period is:

$$\frac{T_{\text{cylinder}}}{T_{\text{sphere}}} = \sqrt{\frac{I_{\text{cylinder}}}{I_{\text{sphere}}}} = \sqrt{\frac{1/2}{2/5}} = \sqrt{\frac{5}{4}}$$

So the new period is:

$$T_{\text{cylinder}} = T_{\text{sphere}} \times \sqrt{5/4}$$

Substituting $T_{\text{sphere}} = 2.50\text{s}$:

$$T_{\text{cylinder}} = (2.50\text{s}) \times \sqrt{5/4} = (2.50\text{s})(1.118) = 2.795\text{s} \approx 2.80\text{s}$$

The new period is 2.80s.
The fractional increase is:

$$\frac{T_{\text{cylinder}} - T_{\text{sphere}}}{T_{\text{sphere}}} = \frac{2.8\text{s} - 2.5\text{s}}{2.5\text{s}} = 0.12 = 12\%$$

The percentage increase is 12% as predicted.

Making Sense of the Answer: *Ratio reasoning is the cleanest tool for problems like this. We did not have to recompute the moment of inertia or re-derive the period. We used the fact that* $T \propto \sqrt{I}$ *and that the only quantity that changed between the two experiments was the moment of inertia. Two lines of algebra and one square root delivered the answer. The same trick works whenever a problem describes 'the same system but with one parameter changed': identify what is unchanged, write a ratio with the unchanged thing cancelling, and substitute.*

Think Like a Physicist: *The 12% difference between the sphere and cylinder periods, for the same mass and radius, is small but measurable. A torsional pendulum experiment that times 100 oscillations of each body can resolve this difference comfortably. Conversely, if you suspended an unknown solid body from the same wire and measured the period, comparing it to what you would expect for a sphere or a cylinder of the same mass and radius would reveal the body's shape category. If the unknown body's period were 2.5s, it would be sphere-like (mass concentrated near the centre); if it were 2.8s, cylinder-like (mass distributed at all radii); if it were larger still, hollow (mass concentrated near the rim). The torsional pendulum is, in this sense, a moment-of-inertia spectrometer.*

We have seen the compound pendulum (gravity along a lever arm) and the torsional pendulum (elastic resistance of a twisted wire). The third and last of the rigid-body oscillators in this chapter is a body that rolls back and forth inside a curved bowl, without slipping.

ROLLING OSCILLATION ON A CURVED SURFACE

We have now seen two of the three rigid-body oscillators of this chapter. The compound pendulum used gravity acting along a lever arm; the torsional pendulum used the elastic resistance of a twisted wire. The third and last is a body that rolls back and forth, without slipping, in a curved bowl. Gravity is again the restoring agent, as in the compound pendulum, but the body's motion now combines two ingredients we have so far kept separate: translation of the centre of mass along an arc, and rotation of the body about its own axis. Rolling links these two motions, in a strict ratio set by the body's radius. The result is a period that is longer than the simple ball-in-bowl period we derived earlier in this chapter, by a factor that depends only on the body's shape, not on its size or its mass.

There is a real reason to do this case carefully. The ball-in-bowl analysis we did earlier treated the bowl as frictionless, so that the ball merely slid up and down the curve without spinning. That idealisation is convenient, but it is rarely realistic. A marble in a steel mixing bowl rolls. A ball-bearing in a curved race rolls. A bottle rocking inside a half-pipe rolls. Rolling is the rule, not the exception, and it adds a measurable amount of inertia to the motion. The aim of this section is to make that addition precise.

The setup and the rolling constraint

A body of radius r (a sphere or a cylinder) is placed inside a smooth bowl whose inner surface is a portion of a sphere of radius R. The geometry is exactly that of the ball-in-bowl figure used earlier in this chapter. We reproduce it here for convenience, but the physics is now different: the body rolls without slipping, instead of sliding.

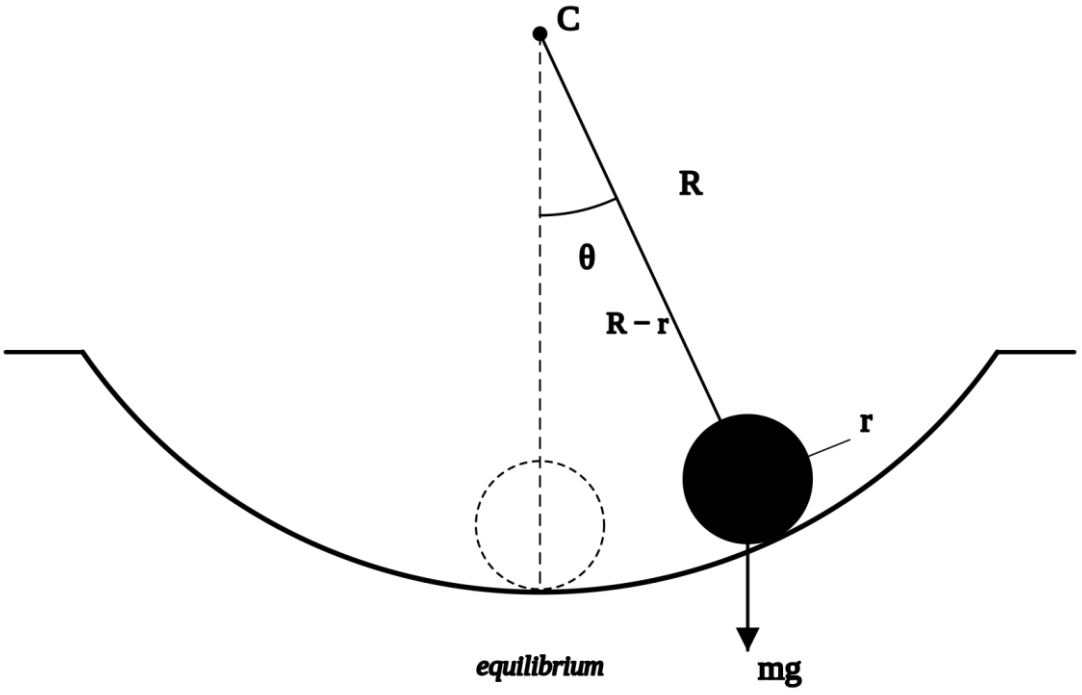


Figure: ***A body of radius r rolling without slipping inside a bowl of radius R****. The centre of the body traces an arc of radius R−r about the bowl's centre of curvature C. The angle θ measures angular position; gravity Mg supplies the restoring torque. Same geometry as the earlier ball-in-bowl figure, but the body now rolls.*

Let θ be the angle, measured at the bowl's centre of curvature C, between the vertical and the line from C to the centre of the rolling body. When the body sits at the bottom of the bowl, θ = 0; when it has rolled a small angle up either side, θ is positive. The centre of the body moves along a circular arc of radius R-r, so its arc length from the bottom is:

$$s = (R-r)\theta$$

Differentiating gives the linear speed of the centre of mass:

$$v = (R-r)\frac{d\theta}{dt}$$

Because the body rolls without slipping, its contact point with the bowl is instantaneously at rest. This forces a strict relation between the linear speed of the centre and the angular speed ω_s of the body about its own centre:

$$v = r\omega_s \quad \Longrightarrow \quad \boldsymbol{\omega_s = \frac{v}{r} = \frac{R-r}{r}\frac{d\theta}{dt}}$$

Notice the amplification: while the line from C sweeps through an angle $d\theta$, the body itself spins through $\frac{R-r}{r}d\theta$, which is much larger when r is small compared with R. A marble in a salad bowl spins many times faster about its own axis than about the bowl's centre. The bowl's curvature does small things; the marble's rotation does big things. Both contribute to the kinetic energy.

Translation plus rotation: the effective inertia

The body's total kinetic energy splits into two terms, one for translation of the centre of mass and one for rotation about it:

$$KE = \frac{1}{2}Mv^2 + \frac{1}{2}I\omega_s^2$$

Here v is the speed of the centre of mass of the rolling body, ω_s is its spin angular speed, and I is its moment of inertia about the centre.

Use the rolling constraint to replace ω_s with v/r:

$$KE = \frac{1}{2}Mv^2 + \frac{1}{2}I\frac{v^2}{r^2} = \frac{1}{2}(M + \frac{I}{r^2})v^2$$

Writing the moment of inertia in terms of the radius of gyration, $I = Mk^2$:

$$KE = \frac{1}{2}M(1 + \frac{k^2}{r^2})v^2$$

The factor in brackets is the effective inertia ratio. The rolling body behaves, for kinetic-energy purposes, like a sliding body of mass $M(1 + k^2/r^2)$. Rolling adds inertia. The hollower the body, the more inertia is added; the more concentrated its mass near its centre, the less. For the four standard cases:

$$\text{solid sphere: } \frac{k^2}{r^2} = \frac{2}{5}; \quad \text{solid cylinder: } \frac{1}{2}; \quad \text{thin spherical shell: } \frac{2}{3}; \quad \text{thin cylindrical shell: } 1$$

So the four effective-mass factors $1 + k^2/r^2$ are 7/5, 3/2, 5/3, 2, a hierarchy that mirrors the rotation-of-rigid-body table you have memorised. These four numbers are all you need for the rest of this section.

Period in the rolling bowl

We now stitch the energy method together, beginning with the potential energy. Looking back at the figure, when the body sits at the bottom of the bowl ($\theta = 0$), its centre is at the lowest point of its arc, a vertical distance R-r below C. When the body is displaced by an angle θ, its centre is at vertical distance $(R-r)\cos\theta$ below C. The body has therefore risen above its equilibrium height by:

$$h = (R-r) - (R-r)\cos\theta = (R-r)(1 - \cos\theta)$$

Multiplying by Mg gives the gravitational potential energy:

$$PE = Mgh = Mg(R-r)(1 - \cos\theta)$$

For small θ, $1 - \cos\theta \approx \theta^2/2$ (the small-angle approximation, exactly as in the simple and compound pendulum sections). So:

$$PE = \frac{1}{2}Mg(R-r)\theta^2$$

Now the kinetic energy, with $v = (R-r)\frac{d\theta}{dt}$, from the equation we derived earlier:

$$KE = \frac{1}{2}M(R-r)^2\left(1 + \frac{k^2}{r^2}\right)\left(\frac{d\theta}{dt}\right)^2$$

Total energy is conserved (constant) during the motion (no friction does any work, since the rolling-contact friction acts at a point that is instantaneously at rest). So the time derivative of E = KE + PE is zero. So:

Differentiating each term:

$$\frac{d}{dt}\left(\frac{d\theta}{dt}\right)^2 = 2\frac{d\theta}{dt}\frac{d^2\theta}{dt^2}, \qquad \frac{d}{dt}\theta^2 = 2\theta\frac{d\theta}{dt}$$

Therefore:

$$\frac{dE}{dt} = M(R-r)^2\left(1 + \frac{k^2}{r^2}\right)\frac{d\theta}{dt}\frac{d^2\theta}{dt^2} + Mg(R-r)\theta\frac{d\theta}{dt} = 0$$

Every term on the left contains a common factor of $d\theta/dt$. Factoring it out:

$$\frac{d\theta}{dt}\left[M(R-r)^2\left(1 + \frac{k^2}{r^2}\right)\frac{d^2\theta}{dt^2} + Mg(R-r)\theta\right] = 0$$

The body is generally moving, so $d\theta/dt$ is not always zero. The bracket must be zero on its own:

$$M(R-r)^2\left(1+\frac{k^2}{r^2}\right)\frac{d^2\theta}{dt^2}+Mg(R-r)\theta=0$$

Now divide both sides by the coefficient of $d^2\theta/dt^2$ (that is, by $M(R-r)^2(1+k^2/r^2)$). The factor M cancels and one factor of (R-r) cancels between the two terms, leaving:

$$\frac{d^2\theta}{dt^2}=-\frac{g}{(R-r)\left(1+k^2/r^2\right)}\theta$$

This is SHM with:

$$\omega^2=\frac{g}{(R-r)\left(1+k^2/r^2\right)},\qquad \mathbf{T=2\pi\sqrt{\frac{(R-r)\left(1+k^2/r^2\right)}{g}}}$$

Compare with the sliding result obtained earlier in this chapter, $T_{slide}=2\pi\sqrt{(R-r)/g}$. Rolling lengthens the period by exactly the square root of the effective-mass factor:

$$\mathbf{T_{roll}=T_{slide}\sqrt{1+\frac{k^2}{r^2}}}$$

A clean, parameter-free amplification. The bowl, the gravity, the body's mass, and even the body's radius all drop out of the ratio. Only the body's shape (through k^2/r^2) survives. This is the rolling fingerprint of the body.

With this understanding, we are now ready to see these ideas come alive through carefully chosen worked examples.

BINDER Example 34

A solid steel ball of radius 2.50cm is placed inside a hemispherical bowl whose inner radius is 30.0cm. The ball is given a small initial displacement and released, and rolls back and forth without slipping. Calculate the period of small oscillations. Take $g = 9.8m/s^2$.

Solution

The ball is a solid sphere, for which $k^2/r^2 = 2/5$, so the effective-mass factor is $1 + 2/5 = 7/5$. The period from the rolling-bowl formula is therefore:

$$\mathbf{T}=2\pi\sqrt{\frac{(R-r)\left(1+k^2/r^2\right)}{g}}=\mathbf{2\pi\sqrt{\frac{7(R-r)}{5g}}}$$

This is a clean, small-letter formula worth memorising.
Substituting:

$$T=2\pi\sqrt{\frac{7(0.300m-0.0250m)}{5\times 9.8m/s^2}}=1.245s\approx 1.25s$$

Making Sense of the Answer: *About 1.25s per full oscillation. If we had ignored rolling and used the sliding formula $T_{slide}=2\pi\sqrt{(R-r)/g}$, we would have got $2\pi\sqrt{0.275/9.8}=1.052s$, about 18% smaller. Eighteen per cent is far above any conceivable measurement error in a laboratory stopwatch experiment, so the rolling correction is not optional.*

Think Like a Physicist: *There is no friction in our final equation, even though we needed friction to make the ball roll in the first place. This is not a contradiction. The static friction at the rolling contact does no work; its point of application is instantaneously at rest, by definition of rolling without slipping. It transmits force between bowl and ball but does not dissipate energy. The energy method, which counts only quantities that change during the motion, is therefore exact. If kinetic friction were present (ball was actually slipping a little), some energy would be lost each cycle and the amplitude would decrease, and that would be the damping problem of the next section.*

REAL Example 35

Mama Clyfar is preparing ugali in the kitchen. While the maize flour is steaming, her son **Arvion** drops a glass marble into a stainless-steel mixing bowl that has been emptied for washing. The bowl has an inner radius of about 18cm and the marble has a radius of about 1.5cm. The marble rolls back and forth across the bottom of the bowl, making a soft clinking sound that fascinates the boy. Estimate the period of the marble's motion, treating the marble as a solid sphere. Then estimate how many oscillations the boy could count in ten seconds before friction noticeably reduced the amplitude.

Solution

We have R = 0.18m, r = 0.015m, and the marble is a solid sphere with $1 + k^2/r^2 = 7/5$. The period is:

$$T = 2\pi\sqrt{\frac{7(R-r)}{5g}} = 2\pi\sqrt{\frac{7(0.18m - 0.015m)}{5 \times 9.8m/s^2}} = 0.965s \approx 1.0s$$

About one cycle per second. In ten seconds, before damping takes noticeable effect, the boy should count roughly **ten oscillations**. (In practice, the marble's amplitude does decrease, because some kinetic energy is lost to imperfect rolling, vibration of the bowl, and air resistance. But for the first ten seconds the loss is gentle: amplitudes drop perhaps 20% over that time, while the period stays essentially constant. Damping changes amplitudes; it does almost nothing to periods, as we shall verify rigorously in the next section.)

Making Sense of the Answer: *This is an experiment any reader can repeat. Find a hemispherical metal bowl. Drop a marble. Time ten oscillations with a phone stopwatch. If the bowl is roughly 18 cm in radius and the marble roughly 1.5 cm, you should measure* ~ 9.7s *for ten cycles. If you measure* ~ 8.2s, *the marble was sliding; if you measure* ~ 13s, *you used a different bowl or a different marble than you thought. The formula is precise enough to catch your mistake.*

Think Like a Physicist: *Why does the formula know nothing about the marble's mass? Because gravity scales mass exactly: doubling* M *doubles both the inertia and the restoring weight, and the two cancel; the same cancellation that makes the period of a simple pendulum independent of bob mass. Why does it know nothing about the marble's material (glass, steel, marble proper, plastic)? Because the formula depends only on shape (*k^2/r^2*) and on the geometric ratio* R-r, *not on density. A glass marble and a steel ball of the same radius and shape category will oscillate identically in the same bowl. This material-independence is one of the most beautiful invariances in mechanics; we shall meet it again in damped oscillations and in resonance.*

HOT Example 36

A solid sphere of mass M and radius r is placed in a bowl of inner radius R and rolls without slipping with period T_0. Without changing anything else, the solid sphere is replaced by a thin spherical shell of the same outer radius r and the same total mass M.

(a) Show that the new period is $T_1 = T_0\sqrt{25/21}$. (b) Express the percentage increase in the period.

Solution

(a) The two effective-mass factors come from the moments of inertia. For a solid sphere, $I = (2/5)Mr^2$, giving factor 7/5. For a thin spherical shell, $I = (2/3)Mr^2$, giving factor 1 + 2/3 = 5/3. The ratio of the new to the old period is therefore:

$$\frac{T_1}{T_0} = \frac{2\pi\sqrt{\frac{5(R-r)}{3g}}}{2\pi\sqrt{\frac{7(R-r)}{5g}}} = \sqrt{\frac{5/3}{7/5}} = \sqrt{\frac{25}{21}}$$

$$\mathbf{T_1 = T_0\sqrt{25/21}}$$

(b) Computing the numerical value:

$$\sqrt{25/21} = 1.0911$$

So the percentage increase is:

$$\frac{T_1 - T_0}{T_0} = \frac{T_1}{T_0} - 1 = 1.0911 - 1 = 0.0911 = 9.11\%$$

Making Sense of the Answer: *The fact that the ratio* T_1/T_0 *came out to* $\sqrt{25/21}$, *a number depending only on* k^2/r^2 *for the two bodies, is the strongest possible illustration of the universality we discussed earlier. The bowl could be 5cm or 5m in radius; the ball could be 1 mm or 10 cm; the gravity could be Earth's or the Moon's; the period would always increase by the same 9.11% on switching from solid sphere to thin shell.*

Think Like a Physicist: *The ratio* $\sqrt{25/21} = 1.0911$ *has a curious quality: it is not a recognisable square root. This is fine for physics, but it warns the student that the universe does not always obey clean Pythagorean ratios. Pick the right reference state, and many of the messy numbers become clean. Pick a different reference, and they need not.*

We have now completed the three rigid-body oscillators of this chapter: the compound pendulum, the torsional pendulum, and the rolling oscillator. Each had its own restoring agent (gravity along a lever arm; an elastic twisted wire; gravity through a curved arc, supplemented by rotational inertia), and each gave rise to a clean period formula of the SHM kind. In every case, the body oscillated forever: undamped, lossless, eternal. Real oscillators do not behave like that. The next section confronts that universal complication.

DAMPED OSCILLATIONS

When we built the equation $a = -\omega^2 x$ at the start of this chapter, we took for granted a quiet and very generous assumption: the oscillator we were studying lost no energy to anything. The mass on the spring would oscillate at the same amplitude on its tenth cycle, on its hundredth cycle, and on its millionth. The pendulum, displaced once at noon, would still be sweeping out the same arc at midnight. The U-tube, the floating cylinder, and the ball in the bowl would all be performing tirelessly at their original amplitudes years after we had walked away from them.

This is, of course, a lie.

You have known it was a lie since Kipanga first stopped pushing the swing. The moment Kipute stopped receiving energy from him, the moment her swing was left to its own physics, it began to die. Each pass through the lowest point was a little less ambitious than the last. The metal chain creaked a little less audibly. The arc shrank. After a few minutes the swing was once again hanging straight down, at rest, looking innocent. Whatever motion it had borrowed from Kipanga had been quietly returned to the air, the bearings, and the chain, like a guest tidying up before leaving.

This section is the proper account of what happened. Real oscillators lose energy. We call the loss of energy from an oscillating system **damping**, and we call the resulting motion a **damped oscillation.** The chapter so far has dealt with the idealised, undamped, eternal version. This section brings the chapter back to the world.

We will keep the discussion qualitative. The mathematics of damped oscillation, with its exponential envelopes and its damping coefficient, belongs to a more advanced course. The physics that you must understand for examination is entirely descriptive, and entirely accessible. What we need are clear definitions, three named cases, one good diagram, and the discipline to read a graph correctly. That is a fair half-day's work.

What damping is

In any real oscillator the body that is moving rubs against something, pushes air out of the way, drags through a fluid, or loses energy to internal friction in the spring or wire that supports it. Each of these processes drains a little of the oscillation's mechanical energy and converts it, irreversibly, into the disordered thermal energy of the surrounding medium. The energy is not destroyed (Chapter 4 told us no energy ever is), but it is no longer available to the oscillation. It has migrated into the random motion of air molecules, into a faintly warmer chain at the top of the swing, into a slightly higher temperature in the metal lattice of the spring. The oscillator does not know where its energy went. It only knows it has less, and it adjusts its amplitude downward in response, the way a household budget adjusts after a long Saturday at Kariakoo.

A **damped oscillation** is therefore an *oscillation in which the amplitude decreases progressively over time because energy is being lost to the surroundings*. An **undamped oscillation** *is one in which there is no such loss, so the amplitude remains constant from cycle to cycle.* Strictly speaking, no oscillation in the real world is fully undamped. What we call undamped is whatever sits close enough to the ideal that its losses are too small to notice over the timescale we care about. The pendulum of a wall clock is very lightly damped, so over the timescale of a single swing its motion appears nearly undamped. However, over longer timescales, energy losses due to air resistance and friction accumulate. Without a continuous energy supply from the clock's mechanism, the pendulum would gradually lose amplitude and eventually come to rest.

A natural question, and a good one to settle now because students frequently get wrong, is *whether damping changes the period of the oscillation as well as the amplitude*. The honest answer has two parts. For **light damping**, which is the only case that interests us at this level and which actually applies to swings, springs, pendulums, and almost every laboratory oscillator, the period is ***approximately* unchanged**. The oscillator goes through its cycle at almost exactly the same rate that it would if no damping were present. What damping does, decisively and visibly, is shrink the amplitude of each successive cycle. If you are asked in an examination whether damping affects the period, the careful answer is *"approximately not, for light damping*". Do not say it has no effect. Say it has approximately no effect, and you will be precise.

What damping definitely does is shrink the amplitude**. That is the thing to remember:** *the amplitude shrinks cycle by cycle, the period barely budges, and the energy that has left the oscillation has gone, by way of friction or air resistance, into the thermal energy of whatever surrounds it.*

Three types of damped oscillation

Different oscillators lose energy at very different rates, and the rate at which energy drains away has a beautiful effect on the kind of motion you observe. There are three qualitatively different cases. You must know all of them, and you must be able to recognise the displacement-time graph of each one on sight. The good news is that the three look completely different on a graph. The bad news is that examiners know this too.

1) Underdamped oscillation

This is the case in which the damping is **light**. Energy leaks out of the system gradually, on a timescale much longer than a single oscillation. The body therefore overshoots its equilibrium position on every pass, swings to the other side, returns, swings back, and oscillates on and on, but with each successive amplitude smaller than the one before. If you sketch the displacement against time you find a sinusoidal oscillation that is being progressively squeezed. Its peaks lie inside a smooth envelope that closes in toward zero. The motion continues until effectively all the energy has been dissipated.

Examples are everywhere. The guitar string dies away after you have plucked it. Kipute's playground swing slows once Kipanga stops pushing. The pendulum oscillates in air. The metal ruler quivers visibly after you have flicked its end against the edge of the desk. In every one of these cases the body oscillates many times before it comes to rest, and each oscillation is a little smaller than the last. Underdamping is what happens when an oscillator is left to its own devices in a world that is only slightly resistive. That is to say, in our world.

2) Critically damped oscillation

This is the case in which the *damping is exactly heavy enough to suppress oscillation altogether, while still allowing the body to return to equilibrium as quickly as possible.* The body does not overshoot, does not swing back the other way, and does not oscillate at all. It simply moves smoothly to its equilibrium position and stays there. Critical damping is the engineer's favourite. It gives the fastest possible return to equilibrium that does not involve oscillation.

There are two examples that you must know. The first is the **shock absorber on a car**. After the car drives over a pothole on the Morogoro road, the suspension must return the body of the car to its resting height as quickly as possible without bouncing. A passenger in a car with underdamped suspension would feel the car oscillating up and down for several seconds after every bump, which would be intolerable on a bad road and dangerous on a good one. Shock absorbers are therefore designed to be critically damped. The second example is the **pointer of a moving-coil galvanometer or ammeter**. When you switch the current on, the pointer must swing across to the correct reading and stop there as quickly as possible, so that the experimenter can read the value without having to wait for the needle to settle. An underdamped pointer would swing past the correct reading several times before coming to rest. An overdamped pointer would creep so slowly toward the correct reading that the experimenter would have given up. Critical damping is the compromise that does the job in the shortest time.

3) Overdamped oscillation

This is the case in which the *damping is so heavy that the body returns to equilibrium more slowly than the critical case.* It still does not overshoot, and it does not oscillate, but it takes longer. The damping force is so large that the body is, in effect, always moving against a thick resistance, and inches its way back to equilibrium without ever quite arriving promptly. Overdamping is rarely what you want when speed is the priority, although it is sometimes what you choose deliberately. A **heavy hydraulic door closer** is overdamped on purpose: the door must close gently and slowly enough that nobody is hit on the way out. The same is true of a **swing door in a hospital ward**, where a rapidly returning door would be a hazard to the patients and a small embarrassment to the architect. **Heavily loaded vehicle suspensions** also tend toward overdamping. The trade-off is a softer ride at the cost of a slower recovery from disturbances.

Three cases, three behaviours: underdamping oscillates with shrinking amplitude, critical damping returns to equilibrium fastest without oscillating, and overdamping returns to equilibrium slowly without oscillating. The same physical system, say, a pointer on a spring, can be made to behave in any of the three ways depending on how much fluid resistance, magnetic braking, or internal friction you arrange around it. This is why an engineer must know all three. The kind of motion is something you choose, not something the system imposes on you. If your engineer has chosen badly, you will know. The market scale will swing, the car will bounce, the needle will dither, and someone will write to the manufacturer.

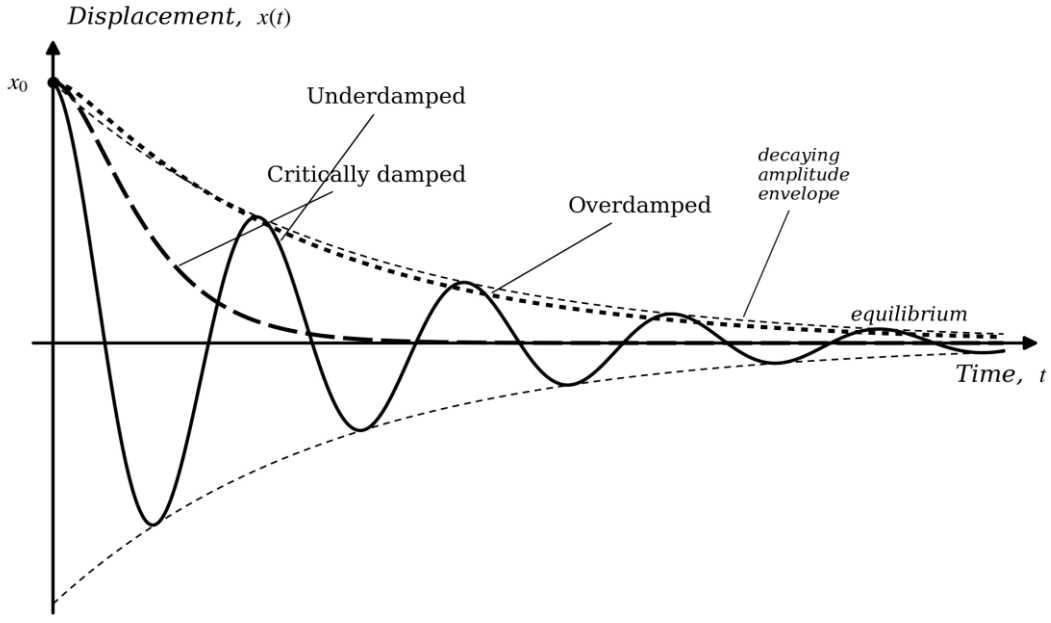


Figure: ***Three types of damped oscillation, all starting from the same initial displacement x_0 at t = 0.*** *The underdamped curve oscillates inside a decaying envelope (dashed); each peak is smaller than the one before. The critically damped curve returns to the equilibrium axis in the shortest possible time without crossing it. The overdamped curve also returns to the equilibrium axis without crossing it, but takes longer than the critically damped curve. Neither the critically damped nor the overdamped curve oscillates. All three curves represent the same kind of physical system; only the strength of the damping has been changed.*

Two worked examples now. They have been waiting patiently while the prose did its work.

BINDER Example 37

(a) Distinguish between damped oscillation and undamped oscillation.

(b) State and describe the three types of damped oscillation.

(c) For each type, sketch the displacement–time graph and give one engineering application.

(d) State whether damping affects the period of oscillation. Justify your answer.

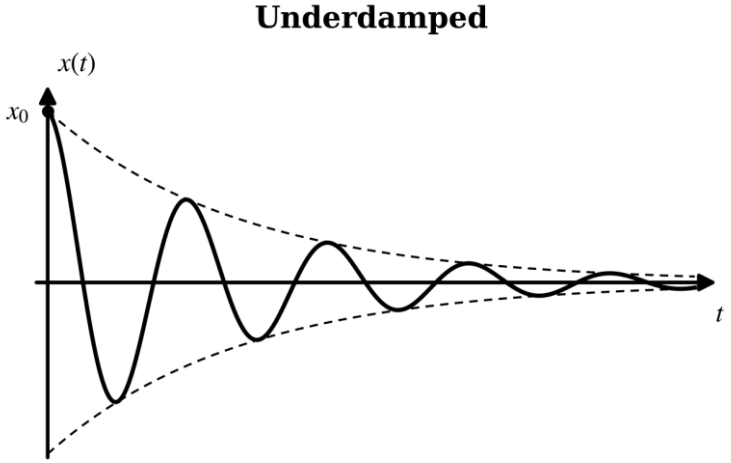
Solution

(a) A damped oscillation is one in which the amplitude decreases progressively over time because energy is lost to the surroundings, typically as thermal energy through friction, air resistance, or viscous drag. An undamped, oscillation is one in which no such energy loss occurs, so the amplitude remains constant from cycle to cycle.

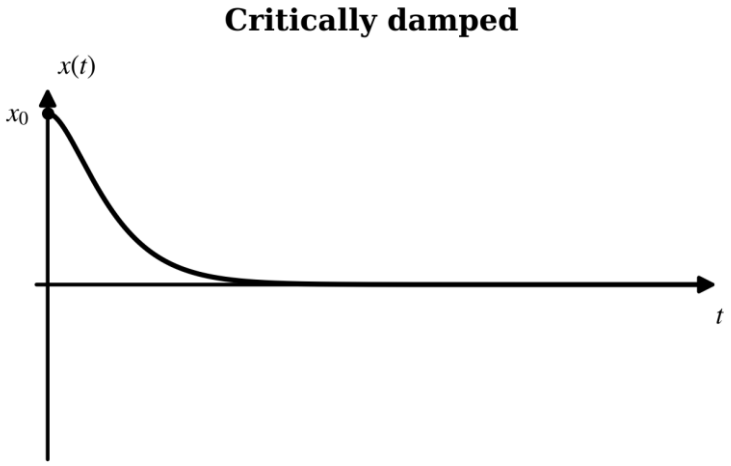
(b) The three types are: (1) underdamped, (2) critically damped, and (3) overdamped.

- In an **underdamped** oscillation the damping is light, so the body oscillates many times before coming to rest, with each amplitude smaller than the last.
- In a **critically damped** oscillation the damping is exactly heavy enough that the body returns to equilibrium in the shortest possible time without oscillating.
- In an **overdamped** oscillation the damping is heavier still, so the body returns to equilibrium without oscillating but takes longer than the critical case.

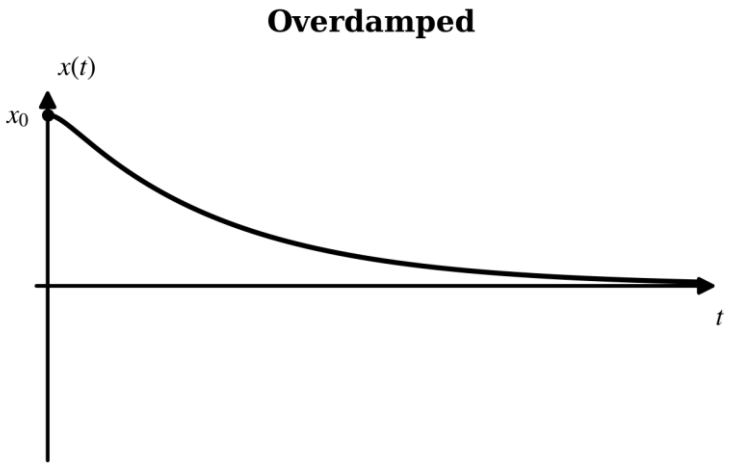
(c) The three displacement–time graphs, are shown below. Each graph starts from the same initial displacement x_0 at $t = 0$, and the equilibrium axis is the horizontal line where $x = 0$.



Underdamped. *Sinusoidal oscillation with progressively decreasing amplitude*. **Engineering example:** *a guitar string* (its vibration is gradually decaying after being plucked).



Critically damped. *Smooth return to equilibrium in the shortest possible time, without crossing the time axis or oscillating.* **Engineering example:** *a car shock absorber (or the pointer of a moving-coil galvanometer).*



Overdamped. *Slow smooth return to equilibrium, again without oscillating.* **Engineering example***: a heavy hydraulic door closer.*

(d) For light damping the period is **approximately unchanged.** Damping decisively reduces the amplitude of each successive cycle, but it does not appreciably alter the rate at which the body completes a cycle. Strictly, the period increases very slightly with damping, but the change is too small to measure for the lightly-damped systems considered at this level.

Making sense of the answer: *Each of the four parts is a small standalone paragraph that addresses exactly what was asked, no more. Examiners reward this structure: a clean (a) / (b) / (c) / (d) layout, a precise distinction where one is asked for, a sketch where one is asked for, and a justified statement where justification is asked for. Part (d) is the most discriminating. Students who write "damping does not affect the period" lose the precision. Students who write "approximately unchanged for light damping" earn it. The word "approximately" is small, but it is doing the work of an entire sentence.*

Think like a physicist: *The two engineering examples for critical damping, the car shock absorber and the galvanometer pointer, are the canonical pair examiners look for. Memorise them as a unit. They both make the same physical demand on the engineer: get to the right answer fast and stop there, with no oscillation. Whether the right answer is the resting height of the car body or the correct value of the current, the engineering principle is the same. Arrive, stop, do not bounce. Politicians could learn from this.*

REAL Example 38

A butcher at Kariakoo market places a 2kg cut of beef on a spring scale. The scale pointer swings back and forth several times before settling on a reading. The butcher is annoyed. He has customers waiting.

(a) Which type of damping does this spring scale exhibit?
(b) Which type of damping would be ideal for this scale, and why?
(c) Explain why each of the other two types of damping would make the scale difficult to use in a busy market.

Solution

(a) The pointer swings back and forth several times before settling, so the scale is **underdamped.**

(b) The ideal scale would be **critically damped.** The pointer would swing across to the reading corresponding to the weight of the beef, stop there at once, and not oscillate. The butcher would read the value the moment the pointer arrived. Critical damping gives the fastest possible return to the equilibrium reading without overshoot, which is exactly what a busy market scale requires.

(c) If the scale were **underdamped,** as in part (a), the pointer would oscillate around the true reading, and the butcher would have to wait several seconds for the swings to die away. With customers waiting, this is too slow. If the scale were **overdamped,** the pointer would creep slowly toward the correct reading and take a long time to arrive. The butcher would still be waiting, only now the pointer would be moving with painful steadiness instead of swinging visibly. Both extremes waste time. Only critical damping, which avoids both the oscillation of underdamping and the sluggishness of overdamping, gives a scale that is fit for a busy market.

Making sense of the answer: *The Kariakoo example shows that the engineering choice between the three types of damping is not abstract. A working scale is a real piece of measuring apparatus, and the wrong damping ruins it for its job. The same underlying physics, the equation of motion of a mass on a spring, produces three completely different commercial outcomes depending on how much resistance the manufacturer designs into the system. The butcher's irritation is, at heart, a complaint about an under-engineered damping coefficient.*

Think like a physicist: *When you are next at Kariakoo, or at any shop that uses a mechanical balance, watch the pointer. If it swings, the manufacturer chose underdamping. If it crawls, the manufacturer chose overdamping. If it arrives at the reading at once and stops dead, the manufacturer chose critical damping, and you are looking at a well-engineered scale. The difference between the three is invisible until you put a 2kg piece of beef on the pan and start watching the pointer. Then it is the only thing in the room. The beef, briefly, becomes philosophy.*

This is the natural place to leave damping. We have learned how a real oscillator quietly loses its energy to the surroundings, what the three named cases of damping look like on a graph, and which engineering choice picks which case. What we have not yet asked is the opposite question. *What happens when energy is being put into the oscillator from the outside, by a periodic driving force?* The answer is the subject of the next section. It contains the explanation we owe Kipanga for the swing on Tuesday afternoon, and the explanation history owes itself for a famous bridge that came down in 1940. Both turn out to be the same explanation.

FREE, FORCED OSCILLATIONS AND RESONANCE

This is the chapter's third visit to the playground swing. The first visit looked at the swing alone, with Kipanga off to one side and Kipute hanging straight down and bored. We solved that swing earlier as a simple pendulum, and we found a single number that told us how long each cycle would take. The second visit, in the section just behind us, watched what happened when nobody pushed. The arc shrank, the chain creaked less audibly each time, and eventually the swing returned to its bored equilibrium. That swing was damped.

The third visit is the one we owe Kipanga an explanation for. *Why did pushing at one rate build the swing higher and higher, while pushing at any other rate seemed to fight it?* The answer is the subject of this section. We meet three new vocabulary words. We meet a graph that holds all of them at once. And we keep an old promise: by the end of the section we will know exactly what went wrong on the bridge at Tacoma Narrows in 1940, and exactly why nobody who knew physics was surprised.

We will keep the discussion qualitative. The full mathematics of forced oscillation is the subject of a more advanced course. What you do need is precise definitions, two pairs of factors that you can list off-hand, and the ability to read a resonance curve correctly. We give you each, in order.

Free oscillations

Imagine a swing displaced once and let go. With no further help from anyone, what does the swing do? It returns through its lowest point, climbs the other side, returns again. Each cycle costs the swing a tiny amount of energy through friction at the pivot and through air resistance, and so each cycle is a little smaller than the one before. Eventually, far in the future, the swing is at rest. While it oscillates, it does so at one particular rate, and the rate is set entirely by the swing itself, by how long the chain is and by how strong gravity is at this latitude. A different swing, with a different chain length, would oscillate at a different rate. The local physical properties of the swing fix the rate, and the rate is the same whether you released the swing from a tiny initial displacement or a generous one.

This rate is called the **natural frequency** of the system. It is written $\mathbf{f_0}$ (sometimes $\boldsymbol{\omega_0}$), with the subscript zero reminding us that nobody is driving the system. For a simple pendulum of length l in a gravitational field of strength g, the natural frequency is $f_0 = (1/2\pi)\sqrt{(g/l)}$. For a mass m on a spring of stiffness k, the natural frequency is $f_0 = (1/2\pi)\sqrt{(k/m)}$. The two formulas are the same shape: a square root that contains the ***restoring agent*** *on top and* the ***inertial agent*** *on the bottom*. We will see this same shape repeated in every oscillator we ever meet.

A system that has been displaced once and is then left alone, with no continuing external driving force, is said to undergo **free oscillation**. "Free" here does not mean unrestrained. It means free of any external driver. Damping is normally still present, so the amplitude does still decay, but no outside source is pumping fresh energy in. The defining feature of a free oscillation is therefore the absence of a driver, not the absence of damping.

A natural question now arises. *What factors determine the period of free oscillation of a mechanical system?* The answer comes straight from the formulas above. There are two factors, no more and no less. The **first** is the *inertia of the system*, mass m for a linear oscillator, or moment of inertia I for a rotational one. The **second** is the *restoring force constant*: spring stiffness k for a mass on a spring, or the equivalent for the system in question (mg/l for a pendulum, 2ρgA for a U-tube, and so on).

One more property of free oscillation deserves a sentence. In the absence of damping, the amplitude of a free oscillation is constant, and equals the initial displacement that set the motion going. In the presence of light damping, the amplitude decays gradually as we saw in the previous section, but the period stays approximately the same. The natural frequency is a property of the system, not of the amplitude or of how violently it was displaced.

Forced oscillations

Now the situation Kipanga walked into. A periodic external force is applied to the system continuously: a hand on a swing, a piston on a spring, a wind on a bridge, a needle on a vibrating phone. The driver chooses the frequency at which the force is

applied, called the **driving frequency**, and this frequency may be quite different from the natural frequency of the system. The driver can push gently or strongly, slowly or quickly, but the system has no say in the rate at which the pushes arrive.

A system that is being driven this way is said to undergo **forced oscillation**. The word "forced" is the right one: the system is not allowed to oscillate at its own natural frequency. After a brief settling time, called the **transient**, the system gives up and oscillates at the driving frequency, not at its own natural frequency. This is a subtle point and a common one to miss. A forced oscillator does not oscillate at the rate it would prefer. It oscillates at the rate the driver dictates. The natural frequency does not disappear. It hides, and waits to matter.

It matters in the size of the response. The amplitude of the resulting forced oscillation depends on two things. The **first** is *how close the driving frequency is to the natural frequency*. The closer the two frequencies, the larger the amplitude of the forced oscillation. The **second** is the *amount of damping in the system*. The lighter the damping, the larger the amplitude.

Resonance

Resonance is the special case of forced oscillation in which the *driving frequency equals the natural frequency of the system*. At this single frequency, energy is transferred from the driver to the oscillator with *maximum efficiency*, and the *amplitude reaches its maximum value*. Every push from the driver arrives at exactly the moment the system is most receptive to it, and so every push adds, rather than partly cancels, the energy already in the system. The amplitude grows. This is what Kipanga discovered on Tuesday afternoon, although he did not yet have the word for it.

Without damping, the amplitude at resonance would grow without limit. Each cycle would add a fixed quantum of energy and lose none, so the amplitude would climb until something physical tore. In practice, damping always exists, and damping cuts the resonance off at a finite amplitude. The lighter the damping, the higher and the narrower the resonance peak. The heavier the damping, the lower and the broader the peak. The graph that shows all of this on a single set of axes is the resonance curve, and you will be expected to sketch it from memory.

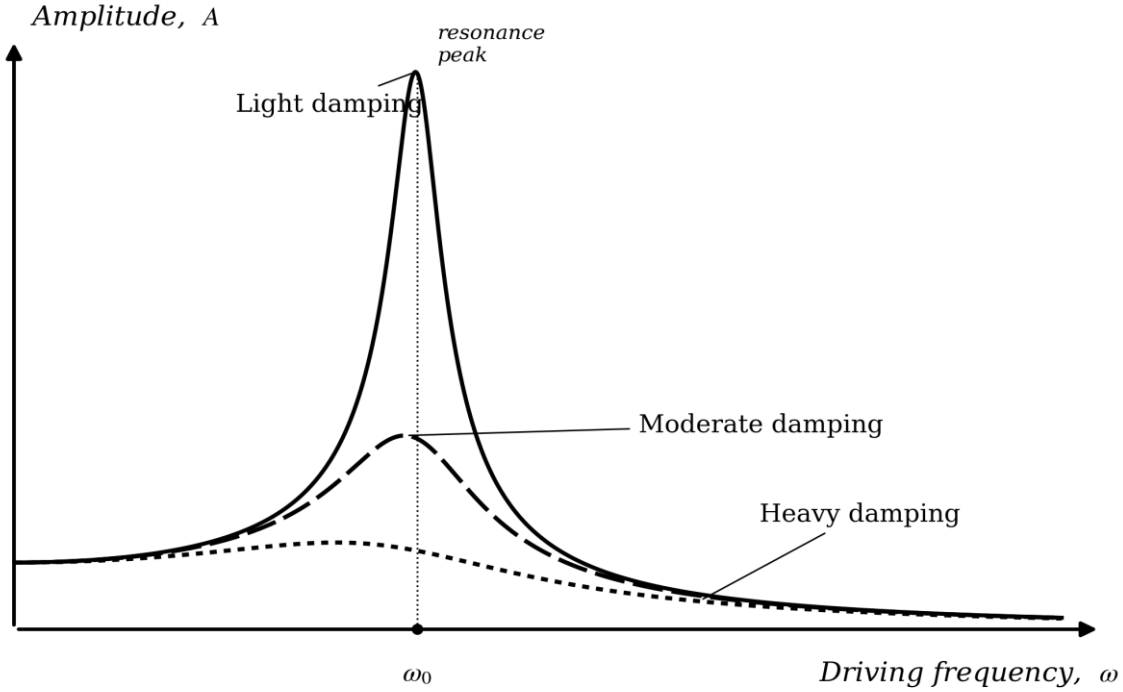


Figure: ***The resonance curve.*** *Amplitude of forced oscillation A is plotted against driving frequency ω, for three levels of damping. With light damping the curve has a tall, narrow peak at the natural frequency ω_0; with moderate damping the peak is shorter and broader; with heavy damping the peak is short, very broad, and may even be slightly displaced to lower frequency. All three curves rise from a small value at very low driving frequencies, peak near ω_0, and fall away at very high driving frequencies.*

Useful resonance

When you want to amplify, resonance is your friend. Engineers and physicists exploit it in at least four places that you will recognise from your own pocket and your own kitchen.

1) Musical instruments

Every instrument with a hollow body, a guitar, a violin, a kora, a marimba, a piano, depends on resonance. The strings (or the bars, or the air column) vibrate at certain natural frequencies, and the body of the instrument is shaped so that it resonates at those same frequencies and amplifies them into something the audience can hear. Without resonance, a plucked guitar string would be barely audible across a quiet room.

2) Tuning a radio

Inside every radio receiver is an electrical oscillator, made of an inductor and a capacitor, whose natural frequency you can adjust by turning the tuning knob (or, in modern devices, by pressing a button). When the natural frequency of this LC circuit matches the broadcast frequency of a particular station, the circuit resonates with the radio waves arriving from that station and ignores the rest. Tuning a radio is, at heart, the engineering of resonance.

3) Microwave ovens

The oven generates electromagnetic waves at a frequency close to one of the natural rotational frequencies of the water molecule. Water molecules in your *sufuria* absorb these waves resonantly, rotate vigorously, and convert the absorbed energy into heat. The food cooks. Most of the rest of your kitchen, lacking water, is left cold.

4) Magnetic resonance imaging (MRI)

The scanner places the patient in a strong magnetic field, which makes hydrogen nuclei in the body precess at a particular frequency. A pulsed radio wave at exactly that frequency causes the nuclei to absorb energy resonantly, and the energy they re-emit a moment later is used to build the image. The scan is, in effect, a controlled resonance experiment performed on every cell of the body at once.

Harmful resonance

When you do not want to amplify, resonance is your enemy. Engineers spend a great deal of design time arranging that natural frequencies do not match the dominant frequencies of any forces the system is likely to experience. The price of forgetting this discipline is steep, as the following examples have, over the years, made clear.

1) Tacoma Narrows Bridge, USA, 1940

A steady wind across the deck of the newly opened bridge happened to drive the deck at a frequency close to one of its torsional modes. The deck twisted with growing amplitude, and on the morning of 7 November 1940 it tore itself apart and collapsed into the strait below. There is film of the collapse. It is one of the most studied engineering failures in history, and it is the reason every modern long-span bridge is designed with explicit attention to its torsional and bending natural frequencies.

2) Buildings in earthquakes

An earthquake is a broad spectrum of seismic frequencies passing through the ground. If a building's natural frequency happens to lie within the dominant band of the earthquake, the building shakes resonantly with growing amplitude until it cracks or falls. Modern earthquake-resistant design tunes a building's natural frequencies far away from the typical earthquake spectrum, and adds heavy dampers (often huge masses suspended on springs near the top of the building) to absorb energy that resonance would otherwise turn into damage.

3) Soldiers crossing footbridges

Marching soldiers walking in unison apply periodic force to the bridge at the rhythm of their step. If that rhythm matches one of the bridge's natural frequencies, the cumulative effect can be enough to set the bridge oscillating dangerously. This is why soldiers are ordered to break step when crossing any footbridge. The order looks superstitious; the physics behind it is not.

4) Engine vibration in vehicles

Every engine produces vibrations at characteristic frequencies set by its rotation rate. Certain panels, exhaust pipes, mirrors, or trim pieces in a poorly designed vehicle have natural frequencies that match those vibrations at common engine speeds. The result is the buzzing rattle that drivers know well, and that returns at exactly the same engine speed every time. The fix, when it can be afforded, is to retune the natural frequencies of the offending parts away from the engine's working range.

The same physics drives the swing higher and brings the bridge down. Resonance is indifferent to consequences. It does what the equations say it will do, and the engineer who knows the equations decides whether resonance helps or hurts. We carry that thought into the worked examples, and into the conversation Kipanga has been waiting all chapter to have.

Two worked examples now. The first is examination practice, dressed in the familiar four-part clothing. The second is a conversation.

BINDER Example 39

(a) Distinguish between free oscillation and forced oscillation.

(b) State two factors that determine the period of free oscillation of a mechanical system.
(c) State two factors that determine the amplitude of a forced oscillation.
(d) Give one example of resonance that is useful in everyday life and one example that is harmful. Explain the physics in one sentence each.

Solution

(a) A **free oscillation** is one in which the system is displaced once and then left to oscillate under its own restoring force only, with no continuing external driving force. The system oscillates at its own natural frequency. A **forced oscillation** is one in which a periodic external driving force is continuously applied to the system, and the system oscillates at the driving frequency, not at its own natural frequency.

(b) The two factors are:

- the **inertia** of the system (mass m for a linear oscillator, or moment of inertia I for a rotational oscillator), and
- the **restoring force constant** of the system (spring stiffness k, or the equivalent restoring parameter for the system in question).

(c) The two factors are:

- how **close the driving frequency is to the natural frequency** of the system (the closer they are, the larger the amplitude), and
- the **amount of damping** present (the lighter the damping, the larger the amplitude).

(d) **A useful example:** tuning a radio. The natural frequency of an LC circuit inside the receiver is adjusted to match the broadcast frequency of the desired station, so that the circuit resonates with the radio waves from that station and amplifies them while ignoring others.

A harmful example: the Tacoma Narrows Bridge collapse of 1940. A steady wind drove a torsional mode of the bridge deck at its natural frequency; the resonant amplitude grew until the deck tore itself apart and the bridge fell into the water below.

Making sense of the answer: *The four parts test the four most frequently examined ideas of this section: the free / forced distinction, the two factors fixing the natural frequency, the two factors fixing the amplitude of the response, and one beneficial and one harmful example of resonance. The structure of the answer mirrors the structure of the question. Examiners reward this. Notice that parts (b) and (c) are answered with two named factors each, no fewer and no more, and that each factor is given in the precise wording the section uses. The candidate who writes "the mass and the springiness" will earn fewer marks than the candidate who writes "the inertia of the system and the restoring force constant". Vague paraphrases lose marks; the section's own phrasing earns them.*

Think like a physicist: *Notice how part (d) has no calculation. Resonance is a phenomenon that can be tested, and is regularly tested, with no numbers at all. The grade depends entirely on the precision of the example and the precision of the one-sentence explanation. Pick examples whose physics you can explain in one breath. The radio and the Tacoma Narrows Bridge are good for this: each has a one-sentence story, and each is famous enough that the examiner will recognise it without needing context. Resist the temptation to invent exotic examples that you cannot explain. The familiar example with a clean explanation outscores the exotic example with a vague one, every time.*

REAL Example 40

It is the same school field, two days later. Kipanga finds Mr. Akilikubwa on the bench by the boundary wall. Kipute is sitting near him, looking impatient with both of them in equal measure. The afternoon is hot. The swing, in the distance, is empty.

Kipanga: *"Sir, about the swing on Tuesday. I think I have understood."*

Mr. Akilikubwa: *"Tell me."*

Kipanga: *"For the first few minutes, every push I gave was undoing what the swing was already doing. I was pushing while it was rising, then pushing again while it was on its way back. The pushes cancelled. The swing barely grew. Then I waited. I pushed at the moment Kipute crossed the lowest point and was moving away from me. Then I pushed again, at exactly the same moment of the next cycle. And again. The pushes added up. The swing rose."*

Mr. Akilikubwa: *"And what is the name of that matching?"*

Kipanga: *"Resonance."*

Kipute: *"It took him three lessons to remember that word."*

Mr. Akilikubwa: *"Be patient with him, Kipute. He had to feel it before he could name it. That is the right order, when there is no examiner standing over you."*

Kipanga: *"But why did the swing build so high? At one point I thought it would loop right over the top bar."*

Mr. Akilikubwa: *"Because at resonance the energy you put in adds to the energy already in the system, every cycle, with no cancellation. The amplitude grows for as long as you keep pushing at the right frequency, until something stops it. In your case the swing eventually rose so high that the small-angle approximation broke down. The simple-pendulum equations stopped applying. Kipute became briefly airborne."*

Kipute: *"I forgive him. The physics was beautiful."*

Mr. Akilikubwa: *"And there is your last lesson. The same resonance that built that swing has, at other times, built more dangerous things. A wine glass cracks when a singer holds the right note. An earthquake destroys a building whose natural frequency happens to match the dominant tremor. A wind plays the wrong tune across a bridge at the wrong speed, and the bridge accepts the invitation."*

Kipanga: *"Tacoma Narrows."*

Mr. Akilikubwa: *"Tacoma Narrows. November 1940. The bridge was lightly damped, the wind was steady, and the wind happened to drive a torsional mode of the bridge deck at exactly the natural frequency of that mode. The deck twisted with growing amplitude until it tore itself apart. There is a film. You have seen it."*

Kipanga: *"Yes, sir."*

Mr. Akilikubwa: *"The wind did not know it was tuning a bridge. The bridge did not know it was singing back. Resonance does not know what it is doing. It is indifferent to consequences. It does the same thing for a swing, for a glass, for a bridge, for the cavity behind the soundboard of a violin. The mathematics is the same. The outcome depends only on what we choose to put in resonance with what."*

Kipute: *"Resonance is indifferent. It does exactly what the physics says."*

Mr. Akilikubwa: *"Exactly. We use it where we want to amplify, and we avoid it where we want to survive. The tuner of a violin and the engineer of a bridge are working on opposite sides of the same equation."*

Mr. Akilikubwa: *"The bridge did not know it was falling."*

The bridge did not know it was falling. The string, as it happens, knows even less. Mr. Akilikubwa's image of the violin and the bridge working on opposite sides of the same equation has one more thing to give us, and we collect it in the next section. So far the oscillations of this chapter have all sat in one place: the swing returned to its bench, the spring returned to its natural length, the bridge twisted around its own deck. None of them travelled. In the next section the oscillations begin to move. We pull a string tight between two fixed points, pluck it, and watch a transverse disturbance run along its length at a perfectly definite speed, set by two perfectly ordinary quantities: how tightly the wire is stretched, and how heavy a metre of it is. The speed has a formula. The formula has a derivation. And the formula is the doorway from oscillation, which has occupied this chapter, into the physics of waves, which will occupy the next.

VELOCITY OF VIBRATION OF A STRING

We are about to close the chapter with a question that has waited politely all the way through it. When you pluck a stretched string, *what determines the speed at which the disturbance runs along the string?* Mr. Akilikubwa's parting line about the violin and the bridge working on opposite sides of the same equation was a hint. We are now in the workshop of the violin, the *zeze*, and the steel guitar, and we ask the question that every instrument maker has had to answer in the dark, by ear: *how do I make this string vibrate at a particular pitch?*

The answer comes in two halves. The first half, given here, is the speed at which a transverse vibration travels along the string. The second half, given in the next chapter, combines that speed with the boundary condition that the two ends of the string are clamped, and from that combination it pulls out the harmonics, the pitches, and the entire physics of stringed instruments. This section is the doorway. The next chapter is the room behind it.

The section also serves a methodological purpose. We do not derive the wave-on-a-string speed by writing down a wave equation and solving it. That belongs to a more advanced course. We derive it by dimensional analysis, a method that uses nothing but the units of physical quantities to deduce the form of an equation. Dimensional analysis cannot tell you everything. It cannot fix numerical constants. But it can fix the shape of an equation, and for many physical questions the shape is what you need. Learning the method here is, on its own, worth the section.

What the speed could possibly depend on

Before writing any equation, ask physically: what properties of the string could possibly determine *how fast a transverse disturbance travels along it?* Reach for two answers, both of which the body knows before the algebra arrives.

The first is the **tension T** in the string. A more tightly stretched string carries vibrations faster, the way a tightly stretched bedsheet snaps back into place faster than a slack one. Tension provides the restoring force that pulls a displaced segment back toward equilibrium, and a stronger restoring force gives a faster wave.

The second is the **linear mass density μ,** the mass per unit length of the string. A heavier string carries vibrations more slowly. More inertia per metre, more sluggish response. The Greek letter μ is the standard notation, and $\mu = m/L$ where m is the mass of a length L of string.

These two physical arguments are doing real work. They identify, on physical grounds alone, the only two quantities that could plausibly enter the answer. Resist the temptation to add other quantities. The wave speed does not depend on the frequency of vibration, on the amplitude of the wave, on the shape of the wave, on how the wave was launched, on the colour of the string, or on the brand of the guitar. It depends only on the medium that the wave is travelling through, which here means only on tension and linear mass density. A wave's speed is a property of the medium. Hold on to this sentence; it is one of the most important sentences in the whole of wave physics.

Derivation by dimensional analysis

Suppose we believe, on the physical grounds just argued, that the wave speed v depends only on the tension T and the linear mass density μ. Then we may write, very generally,

$$v = k \times T^a \times \mu^b$$

where k is a dimensionless constant of proportionality, and a and b are unknown powers that we are about to determine. Demanding that the dimensions of the two sides match will fix a and b uniquely. The numerical constant k cannot be fixed by dimensions alone, and we will return to that limitation honestly at the end.

Write down the dimensions of each quantity

Speed has dimensions of length over time, so $[v] = L\,T^{-1}$. Tension is a force, so $[T] = M\,L\,T^{-2}$. Linear mass density is mass over length, so $[\mu] = M\,L^{-1}$.

Substitute these dimensions into the assumed equation

The dimensionless constant drops out and we are left with:

$$L\,T^{-1} = (M\,L\,T^{-2})^a \times (M\,L^{-1})^b = M^{a+b} \times L^{a-b} \times T^{-2a}$$

Equate the powers of M, L, and T on the two sides

The left-hand side has M^0, L^1, and T^{-1}, so:

$$M:\ a + b = 0 \qquad L:\ a - b = 1 \qquad T:\ -2a = -1$$

From the T equation, $a = 1/2$. From the M equation, $b = -a = -1/2$. The L equation is then $(1/2) - (-1/2) = 1$, which is satisfied automatically. The system is consistent. There is exactly one pair of powers that works.

Substitute back

The wave speed must take the form:

$$v = k \times T^{1/2} \times \mu^{-1/2} = k \times \sqrt{\frac{T}{\mu}}$$

This is as far as dimensional analysis can take us. The dimensionless constant k is not fixed by the method. To fix it we need either a careful experiment with a stretched wire and a stopwatch, or a more advanced derivation from Newton's second law applied to a small element of the string. Either route gives the same answer. Both yield:

$$k = 1$$

and so the speed of a transverse wave on a stretched string is:

$$\mathbf{v = \sqrt{\frac{T}{\mu}}}$$

Pause for a sentence on the method itself. Dimensional analysis has given us the correct shape of the equation, with the right exponents on T and μ, by doing nothing more than counting masses and lengths and times. It has not given us the constant $k = 1$; that is an experimental result confirmed independently. Honest physicists state both halves: dimensional analysis fixes the shape, experiment fixes the scale. Students who believe dimensional analysis is magic come unstuck on the day they meet a problem whose dimensionless constant is not 1.

v is not f: a warning that students must internalise

The result $v = \sqrt{\frac{T}{\mu}}$ tells you **how fast** a transverse disturbance travels along the string. It does not tell you the frequency of vibration, and it does not tell you the pitch. These are different quantities, and confusing them is the most common error students make in the next chapter.

The wave speed v is set by the tension and the linear mass density, both of which are properties of the string itself. The frequency f is set by the boundary conditions: how long the string is, and where it is clamped. A guitarist does not change the wave speed when she plays a different note; the wave speed of her G string at constant tension is the same whether she is playing a melody or sitting silently. What she changes is the length of string that is free to vibrate, and that change in length

changes the frequency without changing the wave speed. We will derive the precise relationship in the next chapter; for now, hold the two quantities apart in your mind.

In plain words: v is the speed at which the wave travels along the string. **f** is the frequency at which the string vibrates. These are different. Tighter strings vibrate at higher pitch because a faster wave fits more cycles per second into the same length, but that is a Chapter 7 argument. Until then: **v depends on T and μ. f depends on the length and the boundary conditions.** Two equations for two questions.

Three worked examples now. The first is single-step substitution, the second a two-step warm-up that traps the careless candidate at the linear-mass-density line, and the third a three-part scaling problem that asks the equation to do real work.

BINDER Example 41

A steel guitar string has a linear mass density of 0.012kg m^{-1}. It is tuned to a tension of 75N. Calculate the speed of transverse vibrations along the string.

Solution

Apply the result of this section directly:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{75\text{N}}{0.012\text{kgm}^{-1}}} = 79.1\text{ms}^{-1}$$

The speed of transverse waves along the string is approximately 79m/s.

Making sense of the answer: *A guitar string and a wave travelling on it at about 79m/s, or roughly 285km/h, may sound impossible until you remember that the wave does not have to travel far. Even a 65cm string is crossed by such a wave in about 8 milliseconds. The string is plucked, the disturbance reaches the bridge in 8ms, reflects, runs back to the nut in 8ms, reflects again, and is therefore making full round trips at a rate of $1/(2 \times 8 \times 10^{-3}) \approx 60$ cycles per second. That figure is suspiciously close to a low musical pitch. The next chapter will turn this rough calculation into the formula for the harmonics of every stringed instrument ever made.*

Think like a physicist: *Notice the unit-by-unit cancellation. The expression $N / (kgm^{-1})$ is, after substituting $N = kgms^{-2}$, simply m^2s^{-2}. Square-rooting gives ms^{-1}, which is a speed. The dimensions are right, and they were guaranteed to be right because we built the equation by dimensional analysis. This kind of unit check is your safety net for every numerical question. If the units come out wrong, the answer is wrong, even when the numbers look plausible.*

BINDER Example 42

A wire of length 0.60m has a mass of 4.8g and is stretched under a tension of 360N. Find (a) the linear mass density μ, and (b) the speed of transverse vibrations on the wire.

Solution

(a) The wire's mass is given in grams; convert to kilograms before dividing, otherwise the units of the answer will be wrong by a factor of a thousand. $4.8\text{g} = 4.8 \times 10^{-3}$ kg. Then

$$\mu = \frac{m}{L} = \frac{4.8 \times 10^{-3}\text{ kg}}{0.60\text{m}} = 8.0 \times 10^{-3}\text{ kgm}^{-1}$$

(b) Apply $v = \sqrt{\frac{T}{\mu}}$:

$$v = \sqrt{\frac{360\text{N}}{8.0 \times 10^{-3}\text{ kgm}^{-1}}} = \sqrt{45\,000\text{ m}^2\text{s}^{-2}} = 212\text{ms}^{-1}$$

The wave speed on this wire is 212m/s.

Making sense of the answer: *Notice that part (a) was the trap. A candidate who substitutes 4.8g and 0.60m directly into m/L gets $\mu = 8.0$kg/m, a thousand times too large. The square root then gives a wave speed a thousand times too small (about 6.7m/s instead of 212m/s). The mistake is invisible until you check units: kilograms and metres must be paired, grams and metres must not. Convert grams to kilograms before any division, and the rest of the calculation runs through cleanly.*

Think like a physicist: *The wire here is ligther per metre than the guitar string of Example 41 ($8.0 \times 10^{-3}kgm^{-1}$ compared with $1.2 \times 10^{-2}kgm^{-1}$, about two-thirds the mass per unit length), and it is under almost five times the tension. Both effects raise the wave speed: less inertia per metre, more restoring force per unit displacement. The wave speed climbs from 79m/s on the guitar string to 212m/s on this wire, a factor of about 2.7. Take a moment to check that the factor of 2.7 is consistent with the changes in T and μ via the $v \propto \sqrt{T}$ and $v \propto 1/\sqrt{\mu}$ scalings; if it is, you have understood the equation rather than just used it.*

HOT Example 43

A vibrating string of length 0.90m and mass 5.4g carries transverse waves at a speed of $120ms^{-1}$. (a) Find the tension in the string. (b) The tension is now increased by 44%. Find the new wave speed. (c) The string is replaced by a new string of the same length and the same tension as in part (a), but with twice the linear mass density of the original. Find the new wave speed.

Solution

(a) Compute the linear mass density first, again converting grams to kilograms:

$$\mu = \frac{m}{L} = \frac{5.4 \times 10^{-3}\ kg}{0.90m} = 6.0 \times 10^{-3}\ kgm^{-1}$$

Then invert $v = \sqrt{\frac{T}{\mu}}$ to give $T = v^2 \times \mu$:

$$T = (120ms^{-1})^2 \times 6.0 \times 10^{-3}\ kgm^{-1} = 14\,400 \times 6.0 \times 10^{-3}\ N = 86.4N$$

(b) The new tension is $T' = 1.44 \times 86.4N = 124.4N$. The shortcut is to notice that $v \propto \sqrt{T}$, so multiplying T by 1.44 multiplies v by $\sqrt{1.44} = 1.2$:

$$v' = 1.2 \times 120ms^{-1} = 144ms^{-1}$$

Increasing the tension by 44% raises the wave speed only by 20%, not by 44%. The square-root relationship between v and T softens every change.

(c) The new string has the same tension T = 86.4N but a new linear mass density $\mu' = 2 \times 6.0 \times 10^{-3}\ kgm^{-1} = 1.2 \times 10^{-2}\ kgm^{-1}$. Compute directly:

$$v'' = \sqrt{\frac{86.4N}{1.2 \times 10^{-2}\ kgm^{-1}}} = \sqrt{7200\ m^2s^{-2}} = 84.9ms^{-1}$$

Or, equivalently, use $v \propto 1/\sqrt{\mu}$: doubling μ divides v by $\sqrt{2} \approx 1.414$, giving $120 / \sqrt{2} \approx 84.9ms^{-1}$, in agreement with the direct calculation.

Making sense of the answer: *The two scalings $v \propto \sqrt{T}$ and $v \propto 1/\sqrt{\mu}$ are the most useful corollaries of the formula. They tell you, without any new calculation, how the wave speed changes when you change either T or μ. Internalise them. To double v you must quadruple T, or you must reduce μ to one quarter of its previous value. Tightening a string by a small amount has only a small effect on the wave speed; halving the mass per unit length only raises v by a factor of $\sqrt{2}$. Real instrument makers know this in their hands.*

Think like a physicist: *Note the discipline of part (b). The slow route is to compute $T' = 1.44 \times 86.4 = 124.4N$, then μ stays the same, then $v' = \sqrt{124.4 / 0.006} = 144m/s$. The fast route is $v' = \sqrt{1.44} \times v = 1.2 \times 120 = 144m/s$. The two routes give the same answer because they are the same physics seen from two angles. The faster method is preferred under examination conditions because it saves time and reduces the risk of arithmetic errors.*

We have arrived at one of the most quietly powerful equations in the chapter. The wave speed on a stretched string depends on exactly two ordinary quantities, the tension and the linear mass density. Dimensional analysis fixed the shape of the equation, experiment fixed the constant, and a handful of substitutions showed how the equation behaves in the workshop. In the next chapter we will combine this speed with the boundary conditions of a clamped string and read off, in a single equation, the entire physics of every stringed instrument ever made: the harmonic frequencies of the guitar, the violin, and the *zeze*.

Before we get there, however, we owe ourselves a short tour. The chapter has produced a handful of equations and eight worked physical systems, and the world has been quietly using all of this for a long time without our help. In the next section we collect six places where the physics of this chapter has earned its keep, from the wristwatch on the desk to the seismometer on the rift, and from the loudspeaker in the phone to the violin in the orchestra pit.

APPLICATIONS OF SIMPLE HARMONIC MOTION

We have spent this chapter chasing a single equation, $a = -\omega^2 x$, into eight different physical systems. We have seen it govern a mass on a spring, a simple pendulum, a U-tube of water, a floating cylinder, a ball in a bowl, a compound pendulum, a torsion wire, and a wave on a stretched string. We close the chapter with a short tour of six places in the world where the same equation has earned its keep. Each example is a working device. Each one depends on a physicist (or a long line of physicists) understanding the equation we have just spent about eighty pages developing. None of the devices know the equation; all of them obey it.

1. Clocks and timing

The oldest application is the easiest to forget. A simple pendulum has a period that depends only on its length and on the local gravitational field, $T = 2\pi\sqrt{\frac{l}{g}}$. A pendulum of length close to 1m (more precisely 0.994m at standard gravity) has a period of exactly 2s. This is the famous "seconds pendulum". It ticks twice per period, once at each extreme of the swing, so the escapement that advances the gears clicks every 1s. The grandfather clock in your aunt's living room is a direct application of $a = -\omega^2 x$: each tick is a half-cycle of a simple pendulum, and the steady tick that keeps the clock running on time is the isochronous property we proved in the simple-pendulum section. The amplitude of the swing decays slightly between winds; the period does not. The clock keeps time.

Modern timekeeping has moved on, but only to a different oscillator. The quartz crystal in a wristwatch is a tiny fork of crystalline silicon dioxide cut to vibrate, by piezoelectric coupling, at exactly 32 768 Hz (which is 2^{15} cycles per second, a number chosen because dividing by two fifteen times, that is 15 successive divisions by 2 gives one cycle per second). The vibration is simple harmonic to extraordinary precision. A digital circuit counts the cycles and ticks the second hand once every 2^{15} of them. A wristwatch costing a few thousand shillings keeps time to about a second a day, which is a million-to-one accuracy in a piece of consumer electronics. The physics is the same physics that governs Kipute's swing.

2. The seismometer

When the ground beneath a seismometer shakes during an earthquake, you might expect the instrument to shake along with it. The clever trick of a seismometer is to use a heavy mass suspended on a long, soft spring or a long pendulum whose natural frequency is far below the frequency of the seismic waves. The mass is, by inertia, almost stationary. The frame of the instrument moves with the ground. The relative displacement between the stationary mass and the moving frame is what the instrument records. A seismometer is therefore an SHM oscillator deliberately tuned to be a bad responder to the driving frequency, so that the frame's motion (which is the earthquake) shows up cleanly as the displacement we want to measure.

The same physics, with the natural frequency tuned in the opposite direction, is at the heart of the accelerometer in your phone. There the mass is light, the spring is stiff, and the natural frequency is high. The mass moves with the phone, and the relative displacement is small but measurable. Whether the design wants the mass to move with the frame or to lag behind it depends entirely on whether the natural frequency sits above or below the driving frequency, which is the resonance physics of the previous section put to engineering use.

3. Vehicle suspension

Every vehicle suspension is a damped mass-spring oscillator. The mass is the car body. The spring is the steel coil between the wheel and the chassis. The damper is the hydraulic shock absorber that we met in the damping section. After the wheel runs over a pothole on the Morogoro road, the spring stores energy and the body rises; the damper then dissipates that energy into heat as the body returns to its resting height. The engineer's choice of damping coefficient, as we discussed in the damping section, is critical. Underdamped suspension produces the bouncing car that takes seconds to settle. Overdamped suspension produces the harsh ride that transmits every bump straight to the passenger. Critically damped suspension does the job in the shortest time without overshoot, and that is what the design aims for.

The natural frequency of a passenger-car suspension is around 1Hz to 1.5Hz, a deliberate choice. Frequencies in this range are slow enough that the human body can absorb them without discomfort. Higher natural frequencies would feel jarring; lower ones would feel like motion sickness. The suspension is, in other words, tuned to the natural frequencies of the human gut. Engineering at its most considerate.

4. The loudspeaker and the microphone

A loudspeaker is forced oscillation made audible. Inside the speaker is a paper or fabric cone glued to a coil of wire that sits in the field of a permanent magnet. When a current oscillating at audio frequency is sent through the coil, the coil experiences a force that varies with the current, and the cone is driven back and forth at the same frequency. The cone, by pushing the air in front of it, sets the air oscillating at that frequency, and the oscillation propagates outward as a sound wave. Every note you have ever heard from a phone, a radio, or a public-address system was produced by a forced simple-harmonic oscillator turning electrical energy into mechanical displacement and then into a sound wave.

A microphone is the same device run backwards. A diaphragm sits in the path of an arriving sound wave. The pressure variations of the wave force the diaphragm into oscillation, and the diaphragm's motion is converted (by induction in a coil, or by capacitance change, or by piezoelectric coupling) into an electrical signal at the same frequency. Loudspeaker and microphone are mechanical inverses of each other, and the same equation, with the same ω, runs in both directions. A physicist who designs one understands the other free of charge.

5. Medical imaging

Two imaging modalities in any modern hospital depend on simple-harmonic physics. The first is ultrasound. A piezoelectric crystal, driven by an oscillating voltage, vibrates as a forced oscillator at frequencies of a few MHz, well above the range of human hearing. The vibrations propagate as ultrasound waves into the body, reflect off boundaries between tissues of different acoustic properties, and return to the same crystal which now acts as a microphone. The timing of the returns is converted into the depth of the reflecting tissue, and the image is built up. The crystal is the same simple-harmonic oscillator we developed in this chapter, only operating at a frequency three orders of magnitude above what any musical instrument could produce.

The second is magnetic resonance imaging, or MRI. We met the principle in the resonance section: hydrogen nuclei in the body precess, in a strong magnetic field, at a particular frequency that depends linearly on the local field strength. By applying gradient magnetic fields and a radio-frequency pulse tuned to the precession frequency, the scanner causes the protons in a thin slice of the body to resonate. The signal that they re-emit a moment later, decaying in amplitude exactly the way the damped oscillations of this chapter do, is used to build the image. An MRI scan is, at its physical heart, a controlled resonance experiment performed on every cell of the patient's body simultaneously, with the decay constant of the signal carrying the contrast information.

6. Musical instruments

Every musical instrument the world has ever made is an application of simple harmonic motion subject to a boundary condition. The string of a guitar is held fixed at both ends, and the standing-wave constraint forces the string to vibrate only at frequencies for which a whole number of half-wavelengths fits between the two clamped ends. Each of those frequencies is a separate simple-harmonic mode. The lowest is the fundamental, $f_1 = (1/2L)\sqrt{\frac{T}{\mu}}$ using the wave speed $v = \sqrt{\frac{T}{\mu}}$ we derived in the previous section, and the higher modes $f_n = n \times f_1$ are the harmonics that give the instrument its character. We will derive this formula explicitly in the next chapter; for now, accept that every plucked, bowed, or struck string in the world is an assembly of simple-harmonic oscillators superposed on the same string.

Wind instruments work on the same principle but with the air column inside the instrument as the oscillating medium and the open or closed ends of the tube as the boundary conditions. Drums and other percussion instruments use the simple-harmonic modes of a stretched membrane. Even the human voice is an air-column resonator: the vocal tract resonates at frequencies set by the geometry of the throat and mouth, and the voice we hear is a superposition of those resonances on the buzzing of the vocal folds.

Six applications, one equation.

That ends the formal teaching of the chapter. The miscellaneous worked examples that follow do not have the courtesy of treating one section at a time. They will combine ideas across sections, and almost always catching the candidate who has been reading the chapter rather than learning it. Read each one through once before reaching for a pen. Identify which sections it draws on. Then solve.

MISCELLANEOUS WORKED EXAMPLES ON SIMPLE HARMONIC MOTION

Example 44

(a) Explain, in physical terms, why a child on a swing must time pushes to match the swing's natural rhythm to make the amplitude grow, and explain what would happen if the pushes arrived at random intervals.

(b) A particle moves so that its acceleration is given by $a = -k_1 \sin(k_2 x)$, where k_1 and k_2 are positive constants. Show that the motion is approximately SHM for small x and find its angular frequency.

Solution

(a) Pushes timed with the swing's natural rhythm act in the same direction as the swing's instantaneous velocity at the moment of each push. Each push therefore does positive work on the swing, adding energy to it on every cycle, and the amplitude builds up. This is resonance: driving frequency equal to natural frequency.

With pushes at random intervals, some coincide with the swing moving towards the pusher (negative work, energy removed) and others with it moving away (positive work, energy added). On average the energy added per cycle is small, and the amplitude does not grow.

(b) For small x: $\sin(k_2 x) \approx k_2 x$. Hence

$$a \approx -k_1 k_2 x$$

which is of the SHM form $a = -\omega^2 x$ with $\omega^2 = k_1 k_2$, giving

$$\omega = \sqrt{k_1 k_2}$$

The approximation, and therefore the SHM, holds only for small x.

Example 45

(a) Account for the experimental observation that the period of a simple pendulum is independent of the mass of the bob.
(b) An SHM particle has angular frequency 4 rad/s and amplitude 6cm. Find the time taken to move from x = 0 to x = 3cm, and from x = 3cm to x = 6cm. Comment briefly on answers.

Solution

(a) The restoring force on the bob along its arc is $F = -mg\sin\theta \approx -mg\theta$ for small θ. Newton's second law along the arc gives $ma = -mg\theta$, so $a = -g\theta$. The mass m cancels: it appears once in the inertia and once in the gravitational restoring force, since gravity itself acts on the same mass. The acceleration, and therefore the period, depends only on g and the length of the pendulum, not on the bob's mass.
(b) Take $x = A\sin(\omega t)$ with A = 6cm, ω = 4 rad/s.
From x = 0 to x = 3cm = A/2: $\sin(\omega t_1) = 1/2$, so $\omega t_1 = \pi/6$, giving:

$$t_1 = \frac{\pi}{6\omega} = \frac{\pi}{24}\ s = 0.131s$$

From x = 0cm to x = 6cm = A: $\sin(\omega t_2) = 1$, so $\omega t_2 = \pi/2$, giving:

$$t_2 = \frac{\pi}{2\omega} = \frac{\pi}{8}\ s$$

From x = 3cm to x = 6cm:

$$t = t_2 - t_1 = \frac{\pi}{8}\ s - \frac{\pi}{24}\ s = 0.262s$$

The time taken in the first leg (x = 0 to x = 3cm) is 0.131s, and for the second leg (x = 3cm to x = 6cm) is 0.262s.
The second leg takes twice as long as the first, because the particle slows as it approaches the extreme position.

Example 46

(a) Explain why a tuned vehicle suspension uses critical damping rather than light damping for passenger comfort, and what would be the consequence of using light damping.
(b) Two SHM oscillators have the same amplitude A and the same angular frequency ω, but they differ in phase by π/3 radians. At the instant the first is at its maximum displacement, find the displacement and the velocity of the second, expressed in terms of A and ω.

Solution

(a) Critical damping returns the suspension to its rest position in the shortest time without oscillating. After the wheel passes a bump, the body of the car rises and is then brought back to its rest height by the damper without bouncing or overshooting. This gives a smooth, quickly-settled ride.

With light damping, the suspension would be underdamped: the car body would oscillate up and down for several cycles after each bump before settling, producing a sea-sickness ride and reducing tyre contact with the road, which is dangerous at speed.

(b) Take the first oscillator as $x_1 = A\sin(\omega t)$, which is at its maximum (x = A) when $\omega t = \pi/2$. The second leads or lags the first by π/3:

$$x_2 = A\sin(\omega t - \pi/3)$$

At $\omega t = \pi/2$ (maximum displacement for the first):

$$x_2 = A\sin(\pi/2 - \pi/3) = A\sin(\pi/6) = A/2$$

Velocity of the second is $v_2 = dx/dt = A\omega\cos(\omega t - \pi/3)$. At $\omega t = \pi/2$:

$$v_2 = A\omega\cos(\pi/2 - \pi/3) = A\omega\cos(\pi/6) = A\omega \times (\sqrt{3}/2)$$

Hence at that instant the second has displacement A/2 and velocity $(\sqrt{3}/2)A\omega$, moving toward larger displacement.

Example 47

(a) Explain why a particle in SHM has its maximum speed at the equilibrium position and zero speed at the extreme positions.
(b) A 0.20kg mass on a spring executes SHM of amplitude 5cm and period 0.50s. Find the displacement at which the kinetic energy is three times the potential energy.

Solution

(a) The total mechanical energy E = KE + PE is conserved. At the extreme positions, the particle is momentarily at rest before reversing direction, so KE = 0 and all the energy is potential: PE = E. As the particle moves towards equilibrium, the restoring force does positive work, converting PE into KE. At the equilibrium position the displacement is zero, so $PE = (1/2)kx^2 = 0$, and all the energy is kinetic: KE = E. The speed is therefore maximum there.
(b) Total energy: $E = (1/2)kA^2$. At displacement x: $PE = (1/2)kx^2$ and $KE = E - PE = (1/2)k(A^2 - x^2)$.
The condition KE = 3PE gives:

$$(1/2)k(A^2 - x^2) = 3 \times (1/2)kx^2$$

$$A^2 - x^2 = 3x^2$$
$$x^2 = \frac{A^2}{4}$$
$$x = \pm\frac{A}{2} = \pm\frac{5}{2}\text{cm} = \pm 2.5\text{cm}$$

The kinetic energy is three times the potential energy at $x = \pm 2.5$cm from equilibrium.

Example 48

(a) Explain how the period of a simple pendulum changes when it is taken from sea level to the top of Mount Kilimanjaro, identifying which physical quantity is responsible for the change.
(b) A spring of stiffness 400N/m has an elastic limit of 24N. A 1.0kg mass is hung from it and set into vertical SHM. Find the maximum amplitude that can be used without exceeding the elastic limit, and the corresponding total energy of the oscillation. Take $g = 9.8\text{m/s}^2$.

Solution

(a) The period of a simple pendulum is $T = 2\pi\sqrt{\frac{l}{g}}$. The length l of the pendulum is unchanged by relocation. The acceleration due to gravity g, however, decreases with altitude (the pendulum is further from the centre of the Earth). At the summit of Kilimanjaro (about 5.9km above sea level), g is smaller than at sea level by roughly 0.2%. A smaller g makes T larger: the pendulum runs slower at the summit. The change is therefore caused by the variation of g with altitude.
(b) The static extension under the weight of the bob is:

$$x_0 = \frac{mg}{k} = \frac{1.0 \times 9.8}{400}\text{ m} = 0.0245\text{m} = 2.45\text{cm}$$

The corresponding tension in the spring at equilibrium is mg = 9.8N. The elastic limit is reached when the total tension equals 24N, i.e., when the extension equals $\frac{24\text{N}}{400\text{N/m}} = 0.060\text{m} = 6.0\text{cm}$. The maximum further stretch from equilibrium is therefore:

$$A_{max} = 6.0\text{cm} - 2.45\text{cm} = 3.55\text{cm} = 0.0355\text{m}$$

The total energy of the oscillation at this amplitude is

$$E = (1/2)kA_{max}^2 = 0.5 \times 400\text{N/m} \times (0.0355\text{m})^2 = 0.252\text{J}$$

Maximum amplitude is approximately 3.6cm, corresponding total energy is 0.25J.

Example 49

(a) Explain why a building whose lowest natural frequency is close to the dominant frequency of seismic waves can collapse during an earthquake of moderate intensity, while a much shorter or much taller building nearby remains intact.
(b) A 5.0kg block hangs from two identical springs, each of stiffness 200N/m. Find the period of vertical SHM when the springs are connected (i) in parallel, (ii) in series.

Solution

(a) During the earthquake the ground oscillates and drives the base of every nearby building at the same frequency. A building whose own natural frequency matches that driving frequency is in resonance: each cycle of the ground transfers energy efficiently into the building's structural oscillation, and the amplitude grows from cycle to cycle. The cumulative stress in the columns and beams eventually exceeds the elastic limit and the building collapses, even though the ground motion itself was moderate.

A much shorter building has a much higher natural frequency, and a much taller one has a much lower natural frequency. In both cases the driving frequency is far from the natural frequency, the resonance condition is not met, and the response amplitude stays small. The buildings shake but do not fail. This is why earthquake-resistant design tunes a building's lowest natural frequency well away from the dominant earthquake band.
(b) Two springs of stiffness k = 200N/m each.
(i) Parallel: effective stiffness $k_p = 2k = 400\text{N/m}$.

$$T_p = 2\pi\sqrt{\frac{m}{k_p}} = 2\pi\sqrt{\frac{5.0\text{kg}}{400\text{N/m}}} = 0.70\text{s}$$

(ii) Series: effective stiffness $k_s = k/2 = 100\text{N/m}$.

$$T_s = 2\pi\sqrt{\frac{5.0\text{kg}}{100\text{N/m}}} = 1.40\text{s}$$

Example 50

(a) Explain why a marching column of soldiers is ordered to break step when crossing a long footbridge.

(b) A grandfather-clock pendulum is made of brass with linear expansion coefficient $\alpha = 1.9 \times 10^{-5}\ K^{-1}$. The clock keeps perfect time at 18°C. The temperature in Dar es Salaam rises to 32°C. Find, in seconds per day, the time the clock loses.

Solution

(a) Marching soldiers strike the bridge at a steady rhythm, applying a periodic vertical force at the frequency of their step (typically ≈ 1.5Hz to 2Hz). A long footbridge has natural frequencies of vertical oscillation that may lie within this same range. If the marching frequency happens to coincide with one of those natural frequencies, the bridge is driven at resonance: each step adds energy, the amplitude grows, and the bridge can sway dangerously or fail. Breaking step makes the foot impacts random in timing, so the cumulative driving force at any single frequency is small, and resonance is avoided.

(b) The period at temperature T depends on the length: $T = 2\pi\sqrt{\frac{l}{g}}$. A small fractional change of length gives

$$\frac{\Delta T}{T} = \frac{1}{2} \times \frac{\Delta l}{l} = \frac{1}{2}\alpha\Delta\theta$$

Substitute $\alpha = 1.9 \times 10^{-5}\ K^{-1}$ and $\Delta\theta = 32 - 18 = 14K$:

$$\frac{\Delta T}{T} = \frac{1}{2} \times (1.9 \times 10^{-5}) \times 14 = 1.33 \times 10^{-4}$$

This is the fractional increase in the pendulum's period. Each second of true time the clock now records as fewer seconds. Over one day (86400s):

$$\Delta t = 86400s \times 1.33 \times 10^{-4} = 11.5s$$

The clock loses about 11.5s per day in the heat.

Example 51

(a) Explain why a moving-coil galvanometer is designed to be critically damped, while a hydraulic door closer is designed to be overdamped.

(b) A simple pendulum of length 80cm hangs from the ceiling of a lift. Find its period when the lift (i) ascends with acceleration $2.0m/s^2$, (ii) descends with acceleration $2.0m/s^2$, (iii) is in free fall. Take $g = 9.8m/s^2$.

Solution

(a) A galvanometer must give a steady reading as soon as the current is switched on. Critical damping is the unique condition that returns the pointer to its correct deflection in the shortest possible time without overshoot or oscillation, so that the experimenter can read the value at once.

A hydraulic door closer must close the door slowly enough that anyone passing through is not struck. Overdamping returns the door to its closed position smoothly but slowly, with no oscillation. Critical damping would shut the door too fast for safe use; light damping would let it swing back and forth and never settle.

(b) Inside the lift, the bob feels an effective gravitational acceleration g_{eff} set by the lift's acceleration.

(i) Lift ascending at $a = 2.0m/s^2$: $g_{eff} = g + a = 9.8m/s^2 + 2.0m/s^2 = 11.8m/s^2$.

$$T_1 = 2\pi\sqrt{\frac{0.80m}{11.8m/s^2}} = 1.64s$$

(ii) Lift descending at $a = 2.0m/s^2$: $g_{eff} = g - a = 9.8m/s^2 - 2.0m/s^2 = 7.8m/s^2$.

$$T_2 = 2\pi\sqrt{\frac{0.80m}{7.8m/s^2}} = 2.01s$$

(iii) Free fall: a = g downward, so $g_{eff} = 0$. The bob is weightless inside the freely falling lift; the restoring force vanishes; the period is undefined (formally infinite). The pendulum simply does not oscillate.

Example 52

(a) Explain why a microwave oven cooks food efficiently while leaving most of the kitchen cold.

(b) A wooden cylinder of length 25cm and uniform density floats with 18cm submerged in fresh water at 20°C, where $\rho = 1000kg/m^3$. The water is heated to 60°C, where $\rho = 983kg/m^3$. Find the percentage change in the period of vertical oscillation of the cylinder.

Solution

(a) The microwave oven generates electromagnetic waves at a frequency close to one of the natural rotational frequencies of the water molecule. Water molecules in the food absorb these waves resonantly, rotate vigorously, and the motion is converted into thermal energy through molecular collisions. The food, which contains water, heats up rapidly. Most of the rest of the kitchen (ceramic plates, glass turntables, walls) contains very little water, has no rotational mode at the oven's frequency, and so does not absorb the waves resonantly. It stays cool.

(b) For a floating cylinder of cross-section A, the equilibrium submerged depth is $h = m/(\rho A)$, and the period of small vertical oscillation is:

$$T = 2\pi\sqrt{\frac{h}{g}} = 2\pi\sqrt{\frac{m}{\rho A g}}$$

Cross-section A and mass m are unchanged by heating; g is unchanged. Only ρ changes. Hence $T \propto 1/\sqrt{\rho}$, and:

$$\frac{T'}{T} = \sqrt{\frac{\rho}{\rho'}} = \sqrt{\frac{1000\text{kg/m}^3}{983\text{kg/m}^3}} = 1.00862$$

The period increases by about 0.86% when the water is heated from 20°C to 60°C.

Example 53

(a) Explain how a seismometer can record an earthquake without being shaken with the ground that is moving beneath it.
(b) A uniform metre rule pivoted at distance d from its centre is timed and found to make 25 oscillations in 40s. Take $g = 9.8\text{m/s}^2$ and find both possible values of d.

Solution

(a) A seismometer suspends a heavy mass on a long, soft spring (or a long pendulum) whose natural frequency lies far below the frequencies of seismic waves (typically less than 1Hz for a seismometer). When the ground moves, the frame of the instrument moves with it, but the suspended mass cannot follow at the seismic frequency: its natural response time is much longer than one cycle of the ground motion. By inertia the mass remains nearly stationary in space while the frame oscillates around it. The relative displacement between the (stationary) mass and the (moving) frame is recorded mechanically or electronically; that record is the earthquake.

(b) Period of one oscillation: $T = 40\text{s}/25 = 1.60\text{s}$. For a uniform rod of length $L = 1.0\text{m}$, the moment of inertia about the centre is $I_{cm} = ML^2/12$, giving radius of gyration about the centre $k = L/\sqrt{12}$ m and $k^2 = 1/12\ \text{m}^2$.

The period of compound pendulum period is:

$$T = 2\pi\sqrt{\frac{k^2 + d^2}{gd}}$$

Squaring and rearranging gives a quadratic in d:

$$d^2 - \frac{T^2 g}{4\pi^2}d + k^2 = 0$$

Coefficient: $T^2g/(4\pi^2) = (1.60\text{s}^2 \times 9.8\text{m/s}^2)/(4\pi^2) = 0.6355\text{m}$.

$$d^2 - 0.6355d + 0.0833 = 0$$
$$d_1 = 0.451\text{m}, \quad d_2 = 0.185\text{m}$$

Both pivot points lie within the metre rule (0 to 0.5m from the centre on either side), and either choice gives the same period $T = 1.60\text{s}$.

Example 54

(a) Explain why a quartz crystal in a wristwatch keeps time more accurately than a mechanical pendulum clock.
(b) A solid sphere of mass 2.0kg and radius 5.0cm is suspended by a wire and executes torsional oscillation with period 1.5s. The sphere is replaced by a thin uniform ring of the same mass and the same radius, suspended by the same wire about a diameter. Find the new period.

Solution

(a) A quartz crystal is a piezoelectric solid that vibrates at a single, sharply-defined natural frequency set by the crystal's mass, geometry, and elastic stiffness. Because the restoring force in a quartz crystal is interatomic and almost free of friction, the damping is extremely light, the resonance peak is correspondingly sharp, and the natural frequency is essentially insensitive to ordinary changes in temperature, mounting, or driving voltage.

A mechanical pendulum, by contrast, depends on its length and on local g; the length expands and contracts with temperature, the pivot wears, and the air drag varies with humidity and altitude. Consequently, a mechanical pendulum clock is less accurate than a quartz crystal.

(b) Period of a torsional pendulum: $T = 2\pi\sqrt{\frac{I}{c}}$, where I is the moment of inertia of the suspended body about the wire and c is the torsional constant of the wire. The wire is the same in both cases, so $T \propto \sqrt{I}$.

For the solid sphere about a diameter through its centre: $I_{sph} = (2/5)MR^2$. For a thin uniform ring about a diameter: $I_{ring} = (1/2)MR^2$. Hence:

$$\frac{T_{ring}}{T_{sph}} = \sqrt{\frac{I_{ring}}{I_{sph}}} = \sqrt{\frac{(1/2)}{(2/5)}} = \frac{\sqrt{5}}{2}$$

$$T_{ring} = (\sqrt{5}/2) \times 1.5\text{s} = 1.68\text{s}$$

The new period is approximately 1.68s.

Example 55

(a) Explain why a vehicle's exhaust pipe sometimes rattles loudly at one particular engine speed but is silent above and below it.

(b) A steel wire of length 80cm and mass 4.0g is fixed at both ends. A small disturbance launched at one end takes 4.0ms to reach the other end. Find the tension in the wire.

Solution

(a) Each loose part of the exhaust system (a panel, a clamp, a heat shield) has its own natural frequency of mechanical vibration, set by its mass and the stiffness of its mounting. The engine produces vibration at a frequency proportional to its rotational speed. At one specific engine speed, the engine's driving frequency happens to coincide with the natural frequency of the loose part, and the two enter resonance: the part oscillates at large amplitude and rattles audibly. Above or below that speed the driving frequency is detuned from the natural frequency, the response amplitude collapses to a small value, and the rattle disappears.

(b) Wave speed:

$$v = \frac{L}{t} = \frac{0.80\text{m}}{4.0 \times 10^{-3}\ \text{s}} = 200\text{m/s}$$

Linear mass density:

$$\mu = \frac{m}{L} = \frac{4.0 \times 10^{-3}\ \text{kg}}{0.80\text{m}} = 5.0 \times 10^{-3}\ \text{kg/m}$$

From $v = \sqrt{\frac{T}{\mu}}$, solving for tension:

$$T = v^2\mu = (200\text{m/s})^2 \times 5.0 \times 10^{-3}\text{kg/m} = 200\text{N}$$

The tension in the wire is 200N.

Example 56

(a) Explain how a vehicle's shock absorbers protect passengers from road bumps, and identify the type of damping they exhibit.

(b) A thin uniform rod of length 1.2m and mass 0.5kg is pivoted at its end and allowed to swing as a compound pendulum of small amplitude. Take $g = 9.8\text{m/s}^2$ and find (i) its period, (ii) the length of the equivalent simple pendulum.

Solution

(a) A shock absorber is a hydraulic damper connected in parallel with the suspension spring. When the wheel passes over a bump, the spring first stores the impact energy as elastic potential energy; the shock absorber then converts that stored energy into thermal energy of the hydraulic fluid as the spring relaxes. The result is that the body of the vehicle returns to its rest height in the shortest possible time without oscillating. Shock absorbers are designed to operate close to **critical damping**; this is the damping condition that achieves the fastest non-oscillatory return to equilibrium and gives the smoothest ride.

(b) For a uniform rod of length L pivoted at its end, the moment of inertia about the pivot (by the parallel-axis theorem) is

$$I_O = (1/12)ML^2 + M(L/2)^2 = (1/3)ML^2$$

The centre of mass is at distance $d = L/2$ from the pivot. The compound pendulum period is:

$$T = 2\pi\sqrt{\frac{I_O}{Mgd}} = 2\pi\sqrt{\frac{(1/3)ML^2}{Mg(L/2)}} = 2\pi\sqrt{\frac{2L}{3g}}$$

(i) Substitute $L = 1.2\text{m}$, $g = 9.8\text{m/s}^2$:

$$T = 2\pi\sqrt{\frac{2 \times 1.2\text{m}}{3 \times 9.8\text{m/s}^2}} = 1.80\text{s}$$

(ii) Equivalent simple pendulum length L_{eq} has the same period: $T = 2\pi\sqrt{\frac{L_{eq}}{g}}$. Comparing with $T = 2\pi\sqrt{\frac{2L}{3g}}$:

$$L_{eq} = \frac{2}{3}L = \frac{2}{3} \times 1.2\text{m} = 0.80\text{m}$$

Period is approximately 1.80s, equivalent simple pendulum length is 0.80m.

Example 57

(a) Explain why a particle in uniform circular motion can be regarded as the superposition of two SHM oscillations along perpendicular diameters.

(b) A particle moves uniformly around a circle of radius 5cm at 4 rev/s. Find the amplitude, the period, and the maximum velocity of the projection of this motion on a diameter of the circle.

Solution

(a) Let the particle move on a circle of radius A with angular speed ω, and let the centre be the origin. Its position vector is

$$(x, y) = (A\cos(\omega t), A\sin(\omega t))$$

The x-coordinate is $A\cos(\omega t)$, which satisfies $d^2x/dt^2 = -\omega^2 x$: SHM along the x-axis with amplitude A, angular frequency ω. The y-coordinate is $A\sin(\omega t)$, which is SHM along the y-axis with the same A and ω, but a quarter-cycle out of phase with x. The actual motion is the vector sum of these two perpendicular SHMs; the resultant traces out the circle. Hence circular motion is exactly the superposition of two perpendicular SHMs of equal amplitude and equal frequency, ninety degrees out of phase.

(b) The projection of the motion on a diameter is one of the two perpendicular SHMs above.

Amplitude: A = 5cm (the radius of the circle). Period: $T = 1/f = 1/4\text{ s} = 0.25\text{s}$.

Maximum speed of the projection equals the speed of the parent particle around the circle:

$$v_{max} = \omega A = 2\pi f \times A = 2\pi \times 4/\text{s} \times 0.05\text{m} = 1.26\text{m/s}$$

Amplitude 5cm, period 0.25s, maximum velocity 1.26m/s.

Example 58

(a) Explain why the period of free oscillation of a mass on a spring is the same regardless of the size of the initial displacement, while the period of a simple pendulum is independent of amplitude only for small angles.

(b) A glass U-tube of uniform cross-section contains water of total column length 80cm. The water is gently displaced and released. Find the period of the resulting oscillation. Take $g = 9.8\text{m/s}^2$.

Solution

(a) For a mass on a spring, the restoring force is exactly $F = -kx$ for all displacements within the spring's elastic range. The acceleration is therefore exactly $a = -(k/m)x = -\omega^2 x$ with $\omega^2 = k/m$ depending only on the spring stiffness and the mass. Since ω is independent of x, so is the period: amplitude does not enter. The motion is exactly SHM at any amplitude up to the elastic limit.

For a simple pendulum, the restoring force along the arc is $F = -mg\sin\theta$, which is not linear in θ. Only when θ is small can we approximate $\sin\theta \approx \theta$, giving $a \approx -g\theta$ and hence SHM with period $T = 2\pi\sqrt{\frac{l}{g}}$. For larger angles $\sin\theta < \theta$, the restoring acceleration is weaker than the SHM model predicts, and the period grows with amplitude. Isochronism is therefore exact for the spring but only approximate for the pendulum.

(b) For a U-tube of uniform cross-section, displacing the liquid by y on one side raises the other by y, giving a height difference of 2y. The restoring force on the column is $F = -2\rho g A y$; the moving mass is $m = \rho A L$ where L is the total liquid length. Hence

$$a = F/m = -(2g/L)y$$

$$\omega^2 = 2g/L$$

$$T = 2\pi\sqrt{\frac{L}{2g}} = 2\pi\sqrt{\frac{0.80\text{m}}{2 \times 9.8\text{m/s}^2}} = 2\pi \times 0.2020\text{s} = 1.27\text{s}$$

The period of the oscillation is approximately 1.27s.

That ends the chapter's worked-out teaching. What follows is Digging Deeper Exercise 6, the chapter's final hand-off. Pen ready, paper ready, mind clear. The chapter ends when you complete it.

DIGGING DEEPER EXERCISE 6

EXERCISE 6A: BINDER QUESTIONS

Question 1

Kipanga and Kipute argue about the relationship between simple harmonic motion and uniform circular motion. Kipanga claims that a particle moving in a circle of radius A with constant angular speed ω is itself executing simple harmonic motion, since its position varies sinusoidally. Kipute insists that the circular motion is not SHM at all, and that only the projection on a diameter is. Resolve the disagreement, identifying clearly what condition the acceleration must satisfy for motion to be SHM, and showing whether the circular motion meets it.

Question 2

The displacement of a body executing SHM is written $y = A\sin(\omega t + \varepsilon)$. A student claims that the phase constant ε has no physical meaning, since one is free to choose the moment $t = 0$. Identify the error in this claim, and account for what ε actually fixes about the motion once two oscillators with the same A and ω are observed together.

Question 3

Kipanga writes the speed of a particle in SHM as $v = \omega\sqrt{A^2 - y^2}$, and Kipute writes it as $v = A\omega\cos(\omega t + \varepsilon)$. Each insists that the two formulas describe different physical situations. Account for the fact that the two formulas are equivalent, and explain in physical terms why the speed is maximum at $y = 0$ and zero at $y = \pm A$, even though it is at the extremes that the magnitude of the acceleration is greatest.

Question 4

Watching a body in SHM at the moment when its displacement is at the positive extreme, a student reasons: *"the body has stopped at this instant; therefore it has no acceleration at this instant either."* Identify the error in this reasoning, and account for the fact that the body, although momentarily stationary, has acceleration of the largest magnitude it ever has during the motion.

Question 5

Over one complete cycle of SHM, the time-averaged kinetic energy and the time-averaged potential energy of the oscillator are observed to be equal, each being half of the total energy $\frac{1}{2}kA^2$. A student concludes that, since the two averages are equal, the kinetic and potential energies must be equal at every instant. Identify the error in this conclusion, and account in physical terms for the equality of the two averages.

Question 6

A vertical mass-spring system has the same period of SHM, $T = 2\pi\sqrt{\frac{m}{k}}$, whether it is set up at sea level or on top of Mt. Kilimanjaro, where the gravitational field is appreciably weaker. A student concludes from this that gravity plays no part in the motion at all. Identify the error in this conclusion, and account in physical terms for the disappearance of g from the period formula.

Question 7

Kipanga is told that the simple pendulum is "isochronous" and concludes that any pendulum, swung through any amplitude, must have the period $T = 2\pi\sqrt{\frac{l}{g}}$. Kipute insists that this is only an approximation and that the period must in fact grow as the amplitude grows. By referring to the form of the restoring acceleration $a = -g\sin\theta$, and to the approximation by which it is replaced in deriving the formula, account for the fact that Kipute is right and identify the angular range within which Kipanga's claim is acceptable.

Question 8

A uniform metre stick is pivoted at a series of points, the distance h of the pivot from the centre of mass being varied through the experiment. The period of small oscillations is found to fall as h is increased from zero, to reach a minimum, and then to rise again as h approaches the end of the stick. By referring to the period of a compound pendulum,

$$T = 2\pi\sqrt{\frac{k^2 + h^2}{gh}}$$

(where k is the radius of gyration about the centre of mass), account for the appearance of a minimum period.

Question 9

Two everyday systems are observed. The shock absorber on a daladala settles back to rest very quickly after a bump in the **Nyerere** road, with no visible oscillation. A swing door, after being pushed open and released, swings shut, swings half-open, swings shut, and only after several diminishing swings comes to rest. By referring to the three regimes of damped oscillations and to the relative size of the damping coefficient compared with the natural angular frequency ω_0, account for the difference and identify which regime each system is designed to operate in.

EXERCISE 6B: REAL QUESTIONS

Question 10

The shock absorber on a daladala plying the Mbezi-Kariakoo route, when new, returns the vehicle to its rest height after a single bump in the road with no visible oscillation. Six months later, after long use, the suspension is observed to oscillate up and down two or three times after every bump before settling. Identify the regime of damped motion in which the suspension operates in each case (when new and when worn), and account in physical terms for the change.

Question 11

The brass bell at a primary school in Mtwara, when struck once at break time, rings with a clear note that fades steadily over about ten seconds before falling silent. Account in physical terms for the steady fading, and explain why a heavier bell of the same shape and material would ring for longer.

Question 12

A washing machine in a hostel in Dodoma is observed to shake the floor violently when its drum is being driven through one particular spin speed during start-up, but to be smooth at every other speed, including the final one. By referring to the natural frequency of the floor-machine system and to the forcing frequency applied by the rotating drum, account in physical terms for the speed-specific shaking.

Question 13

A driver on the Dar–Morogoro highway notices that her vehicle shakes most uncomfortably at exactly 80km/h, but smoothly at 70km/h and at 90km/h. The road has expansion joints set at a regular interval along its length. By identifying the forcing frequency that the joints impose on the suspension, and by comparing it with the suspension's natural frequency, account in physical terms for the speed-specific shaking.

Question 14

A platoon of police cadets practising drills near the Selander Bridge in Dar es Salaam is told by their instructor to "break step" before crossing. By referring to the forced oscillation of the bridge under a periodic footfall, and to what happens when the footfall frequency matches the bridge's natural frequency, account in physical terms for the safety concern that the instruction addresses.

Question 15

A pendulum clock made by a Dar es Salaam watchmaker keeps perfect time at sea level on the coast. The same clock, carried by a friend up to a hotel on the slope of Mt. Kilimanjaro at about 3000m altitude, is observed to run slow by about one minute per day. By referring to the period formula of the simple pendulum, account in physical terms for the loss of time, and indicate the physical adjustment the watchmaker must make to restore correct timekeeping.

Question 16

The seconds pendulum in Mr. Mushu's physics laboratory has been knocked accidentally and is now keeping time slightly fast. To return it to correct time, the technician adjusts the position of the bob along the rod. State in which direction (up or down along the rod) the bob must be moved, and account in physical terms for the resulting change in period.

Question 17

A spring scale in use by a tomato seller at Makoroboi market in Mwanza shows a steady reading when 2kg of tomatoes is placed on its pan. When the seller jostles the pan, the pointer is observed to oscillate up and down briefly about the steady reading before coming to rest. By referring to the spring-mass system, account in physical terms for the brief oscillation, and identify the period of the pointer's motion in terms of the spring constant of the scale and the mass on the pan.

Question 18

Two mattresses of the same brand in a Kahama guesthouse are observed to differ when a guest sits down: a softer mattress lets the guest oscillate up and down at a slow tempo, with the rise and fall plainly visible, while a firmer mattress of the same outer dimensions barely lets the guest perceive any motion before it settles. By referring to the period of a spring-mass system, account in physical terms for the difference in tempo, and indicate which of the two mattresses has the higher spring constant.

Question 19

A nylon clothesline stretched tight between two posts in a courtyard, when struck with a stick, sends a barely-visible pulse from one post to the other; the same line, when allowed to sag loose, sends a slow visible wave when struck. Account in physical terms for the difference in propagation speed of the two pulses.

Question 20

A heavy mango hanging from a slender branch in Bagamoyo is observed, after the branch is pulled aside and released, to swing back and forth approximately periodically before coming to rest. By referring to the form of the restoring force on the branch-mango system and to the condition for SHM, account in physical terms for why the motion approximates SHM, and identify a specific circumstance under which the approximation would fail.

EXERCISE 6C: HOT QUESTIONS

Question 21

A 2.5kg block hangs in equilibrium from a vertical spring whose natural length is 0.50m and which stretches to a total length of 0.65m under the block. The block is then displaced 0.05m below equilibrium and released from rest. Find the position of the block, relative to the equilibrium position, 1.0s after release, and the maximum kinetic energy of the motion that follows.

Question 22

A simple pendulum has a period of 2.000s at sea level, where $g = 9.80\,\text{m/s}^2$. The same pendulum, when taken to a high-altitude site, is observed to lose 30 seconds per day. The length of the pendulum may be taken to be unchanged. Find the value of g at the new site.

Question 23

A particle executes SHM with amplitude 0.04m and frequency 2.0Hz. Find the displacement at which its kinetic energy equals its potential energy, and the time taken for it to move from $x = 0$ to $x = A/2$, taking the particle to be moving in the positive direction at $x = 0$.

Question 24

A spring of natural length 0.20m stretches to 0.24m when a 0.5kg mass is hung from it in equilibrium. The same spring is then connected end-to-end (in series) with a second, identical spring, and the same 0.5kg mass is hung from the lower end of the combination. Find the period of small vertical oscillations of the mass on the series combination.

Question 25

A pendulum clock has a steel rod that beats seconds ($T = 2.000\text{s}$) at 20°C. The clock is moved to a workshop where the temperature is 35°C. The linear thermal expansion coefficient of steel is $\alpha = 1.2\times10^{-5}\ \text{K}^{-1}$. Find by how many seconds the clock runs fast or slow per day, and state which.

Question 26

A steel wire of length 2.0m, cross-sectional area 0.50mm^2, and Young's modulus $2.0\times10^{11}\ \text{N/m}^2$ hangs vertically from a fixed support. A 4.0kg mass attached to the lower end is given a small downward displacement and released. Find the period of the resulting vertical oscillations.

Question 27

A horizontal nylon string of length 1.50m and total mass 0.030kg is stretched between two posts in a courtyard. A short transverse pulse, generated by a flick at one post, travels to the other post in 0.10s. Find the tension in the string.

Question 28

The displacement of a particle of mass 0.20kg executing SHM is given by $x = 0.05\sin(4\pi t + \pi/3)$, with x in metres and t in seconds. Find the first positive time at which the displacement equals +0.025m, the velocity of the particle at that instant, and the total mechanical energy of the motion.

Question 29

Two simple pendulums, of lengths 1.00m and 1.21m, are released simultaneously from small angles on the same side of the vertical. Find the time at which they next return simultaneously to their initial positions, both moving in the same direction as at the moment of release.

Question 30

A 0.40kg mass on a horizontal spring of stiffness 25N/m executes SHM on a frictionless surface with amplitude 0.10m. Find the kinetic energy of the mass when its displacement from equilibrium is $x = 0.06\text{m}$, and the displacement at which its kinetic energy is twice its potential energy.

Question 31

A boat at a fixed mooring on Lake Victoria moves vertically through a total range of 1.20m (peak to peak), with a period of 4.0s. Treating the motion as SHM, find the maximum vertical acceleration of the boat and the time taken for it to rise from its lowest point to a position 0.30m above the mean position.

Question 32

A uniform wire of length 0.60m and total mass 1.2g is stretched horizontally between two clamps with a tension of 80N. Find the time taken by a transverse pulse to make a round trip from one clamp to the other and back. If the wire is now replaced by a uniform wire of the same length and the same tension but with twice the mass per unit length, find the new round-trip time.

Chapter 7

WAVE MOTION

INTRODUCTION

The morning light on Lake Victoria in Mwanza had the apologetic quality of a guest who arrives too early to a wedding and is not sure where to sit. A bus had deposited the Form Five biology class on a study tour along the lakeshore. The biology teacher had promised fish. There were no fish. There was a great deal of lake, a great deal of mist, and Kipanga, who had been told there would be fish.

After forty minutes of looking at the same patch of water, Kipanga did what bored teenagers have done at lake shores since lakes were invented. He picked up a stone and he dropped it.

The ripples spread outward in widening rings. He watched them go, mildly impressed that a single dropped stone could cause a small geometric event of this much beauty for free. Mr. Akilikubwa, who had not been on the trip officially but had attached himself to it the way physics teachers attach themselves to anything involving water, materialised at Kipanga's elbow. He had with him a small dry leaf which he placed, with the unhurried solemnity of a magician introducing his most famous trick, on the surface of the lake about a metre away from where the ripples were heading.

They both watched. The leaf rose. The leaf fell. The leaf rose. The leaf fell. The ripples passed underneath it and continued on toward the open lake. The leaf, having bobbed, stayed exactly where it had been placed.

Kipanga: *"The leaf didn't move."*

Mr. Akilikubwa: *"Correct."*

Kipanga: *"But the wave moved."*

Mr. Akilikubwa: *"Also correct. Then what exactly moved?"*

Before Kipanga could answer, Mr. Akilikubwa picked up a second stone and threw it further out. A second pattern of ripples bloomed and spread. The two patterns met somewhere in the middle of the spreading region. Where two crests crossed, a taller crest briefly stood up. Where a crest met a trough, the water momentarily went flat, for a single instant, as if the lake had paused to think. Then the two patterns passed through one another, untroubled, untouched, and continued outward. Each ring kept its shape. Each ring kept its speed. Neither one knew, after the meeting, that the other had ever existed.

Mr. Akilikubwa: *"Energy moved outward, Kipanga. Disturbance moved outward. What did NOT move is the water itself. The water molecules merely oscillated up and down, passing the disturbance along to their neighbours like a message whispered down a line of schoolchildren. When the message has gone, the children remain, in their seats, in their original places. That is a wave. And the two ripple patterns added together where they met (displacements summed, not waters) and then walked away from each other unchanged. This is* ***superposition****, and it is the most important principle in the whole of wave physics. We will spend a chapter on it."*

Kipanga: *"All this from one stone?"*

Mr. Akilikubwa: *"All this from one stone. And from a leaf that refused to travel with it. The leaf has just told you, free of charge, that a wave carries energy without carrying matter. The two ripple patterns have just told you that waves can pass through each other and survive. You have witnessed two of the four most important facts about waves before breakfast. The other two are coming. The third is that every point on a wave is doing simple harmonic motion, of the kind discussed in the previous chapter. The fourth is that a wave is a single equation, and that equation is the only equation you need to memorise."*

He picked up a stick and wrote in the wet sand, in handwriting that would have failed an art examination but would have passed a physics one:

$$y(x, t) = A \sin(\omega t - kx)$$

"The whole chapter," he said, *"is about earning that equation. You will earn the right to write it down. You will derive it. You will read parameters off it the way you read a clock. You will use it to predict the speed of seismic waves through Tanzanian rock and the tension in a guitar string and the pitch of the zeze your cousin plays at weddings. You will also, twice in this chapter, watch it become a different equation, a stationary wave, which is what happens when a wave meets its own reflection and has to share the room with it. By the end of these eighty-odd pages you will know why a leaf does not travel with a ripple, why two speakers can make some seats in a hall louder and some quieter, why a guitar string an octave higher is half as long, and why the railway worker on the TAZARA line presses his ear to the rail. All of it from that."*

He pointed at the equation in the sand. A small wave came up the shore, lapped politely at the writing, considered it, and erased it without comment.

Mr. Akilikubwa: *"The lake is also a teacher. It does not believe in keeping notes!"*

Welcome to Chapter 7. The previous six chapters spent themselves on matter, with chapter 6, the most recent, taught you the language of repetition: simple harmonic motion, the rhythm at the bottom of every smooth potential well in the universe.

This chapter is the first to be about the spatial extension of that rhythm. SHM, until now, has been a story about a single point: one mass, one spring, one bob. A wave is what happens when SHM gets shared. One particle is doing simple harmonic motion. Its neighbour is doing the same simple harmonic motion, but a fraction of a second later. Its neighbour's neighbour is doing the same simple harmonic motion, but two fractions of a second later. And so on, all the way along the medium. The result, viewed from above and frozen at one instant, is a sinusoidal pattern in space. Viewed at one fixed point and watched in time, it is a sinusoidal oscillation. Both descriptions are the same wave. Both descriptions live inside the equation Mr. Akilikubwa wrote in the sand.

We also have a promise to keep. At the end of Chapter 6, in the last full section before the chapter ended, we derived the speed of a transverse vibration on a stretched string by dimensional analysis:

$$v = \sqrt{\frac{T}{\mu}}$$

We told you, with a tip of the hat, that the next chapter would do something with that result. Here is the next chapter, and here is the something. We will take that single line for the wave speed, combine it with the requirement that the two ends of a guitar string are tied down so they cannot move, and out will fall the entire physics of stringed instruments: every harmonic, every overtone, every reason a *kora* and a *zeze* can be tuned by ear without an electronic tuner ever leaving its case. The room behind that doorway is this chapter.

And we will need, more often than is comfortable, the language of phase from circular motion in the first book of this book series, ***Foundations of Physics***. Phase is the angular bookkeeping of where a particle is in its cycle, and a wave is built out of particles whose phases differ from each other by amounts proportional to how far apart they are. If the reference circle that produced the SHM equation feels rusty, take ten minutes before the next section and refresh it. The calculations here will not be hard, but they will be relentless about asking you to subtract one phase from another. A student who is fluent in phase finds this chapter delightful. A student who is shaky in phase finds it punishing. The choice is yours and it costs ten minutes.

Two other things to flag before we start. First, this chapter makes a sharp and important distinction between two different velocities. One is the velocity at which the wave pattern travels through the medium. The other is the velocity at which an individual particle moves up and down at right angles to the medium. Confusing the two is the single most common error in this entire chapter. Second, this chapter contains two equations that look almost identical but mean different things, namely:

$$y = A\sin(\omega t - kx) \quad \text{and} \quad y = (2A\cos kx)\sin\omega t$$

The first describes a wave that is travelling. The second describes a wave that is standing still while vibrating in place. Telling them apart is half the chapter's examination work, and we will spend a careful section on it.

That is the chapter. By the time you finish this chapter, you should be able to look at the surface of any pool of water disturbed by anything, and read the physics off it the way a musician reads a score. Let us begin.

WHAT A WAVE IS: ENERGY WITHOUT MATTER

Mr. Akilikubwa's lakeside lecture has done its job. We know, before any equation, two non-trivial facts about a wave: *it carries energy through space, but it does not carry the medium with it.* The water in front of the leaf at the start of the morning is the same water in front of the leaf at the end. The leaf has bobbed politely on the spot while the disturbance has cleared off into the open lake without it. Whatever a wave is, then, it is not a moving object. It is something more interesting: a moving **pattern** of disturbance, hosted by a medium whose particles agree, by the laws of mechanics, to take turns moving up and down so that the pattern can pass.

Here is the formal definition. A **wave** *is a disturbance that propagates through a medium (or, in the case of light, through empty space), transferring energy from one place to another without the bulk transport of the medium itself.* The medium does its work locally, by passing the disturbance along, and then it returns to where it started. The disturbance keeps going. Energy moves. Matter does not.

The point of this section is to make sure that, before we write down the equation of a wave, you have a firm physical picture of three separate things: *what is moving*, *what is not moving*, and *what we mean by every word we are about to use to describe both*. We will start with the two physical kinds of wave that nature offers, then build the small vocabulary that the rest of the chapter will live inside.

Transverse and longitudinal: two ways for a particle to dance

Imagine a long horizontal rope held by two people, one at each end. One person flicks her wrist sharply upward and then back to rest. A pulse runs along the rope toward the other person. The pulse travels horizontally. Each tiny segment of rope, however, has only moved vertically: it has gone up, and then back down, and is now where it was. The **direction of disturbance** (*vertical*) is at *right angles* to the **direction of propagation** (*horizonta*l). A wave with this geometry is called a **transverse wave**. *Water surface waves, the waves on a guitar string,* and *all electromagnetic waves* (light, radio, X-rays) are transverse.

Now imagine instead a long spring. One person pushes one end of the spring sharply along its length and pulls back. A region of compression travels down the spring toward the other person. The pulse travels along the spring. Each coil of the spring, however, has only moved along the spring: it has been squeezed forward into its neighbour and then bounced back. The **direction of disturbance** is *the same as* the **direction of propagation.** A wave with this geometry is a **longitudinal wave**. *Sound waves in air, sound in water,* the waves your eardrum receives at every *Bongo Fleva* song, are all longitudinal. So are *seismic P-waves*, which we will meet again at the end of this chapter when we use them to locate earthquakes.

Two pieces of advice that will pay you back many times in this chapter. First, longitudinal waves are harder to draw than transverse waves, because the disturbance is along the propagation axis and overlaps the picture itself. The standard trick is to draw the longitudinal wave as a sequence of compressions (tightly packed coils) and rarefactions (loosely spaced coils), with a graph of **displacement** of each coil from its rest position drawn separately above the spring. Both pictures describe the same wave. The graph is the trick that lets you compare a longitudinal wave to a transverse one on equal terms. Second, even though sound is longitudinal, when you read its equation it will look exactly like the equation of a transverse wave. The mathematics does not care about the geometry; it cares about the displacement. We will exploit this in the next section to write a single equation that describes both kinds at once.

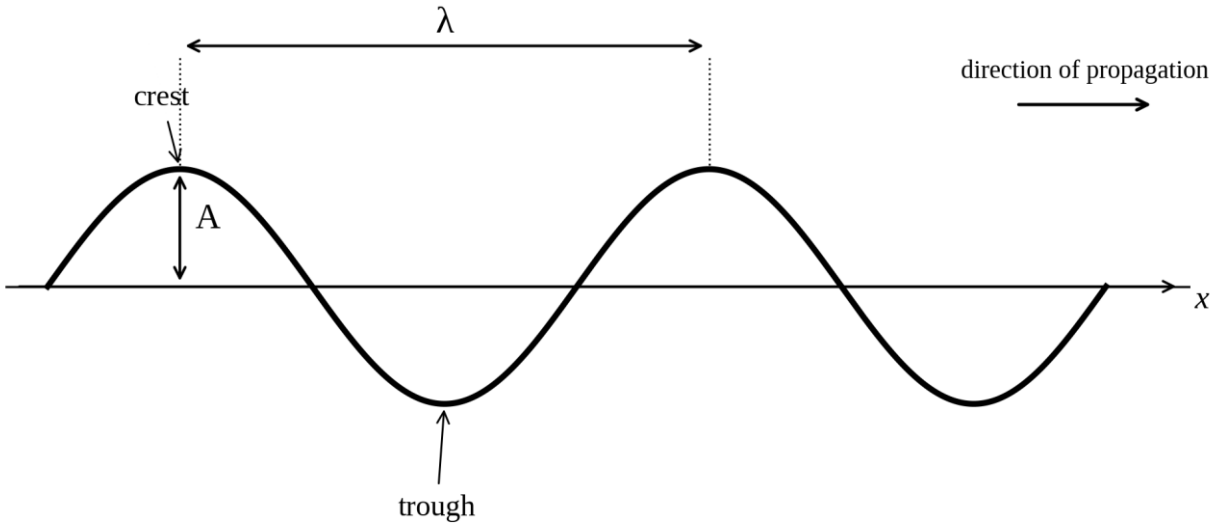


Figure: ***A transverse wave.*** *The medium oscillates vertically (the y direction). The wave propagates horizontally (the x direction). A is the amplitude (peak displacement). λ is the wavelength (distance between consecutive crests, or between any two consecutive points in identical phase).*

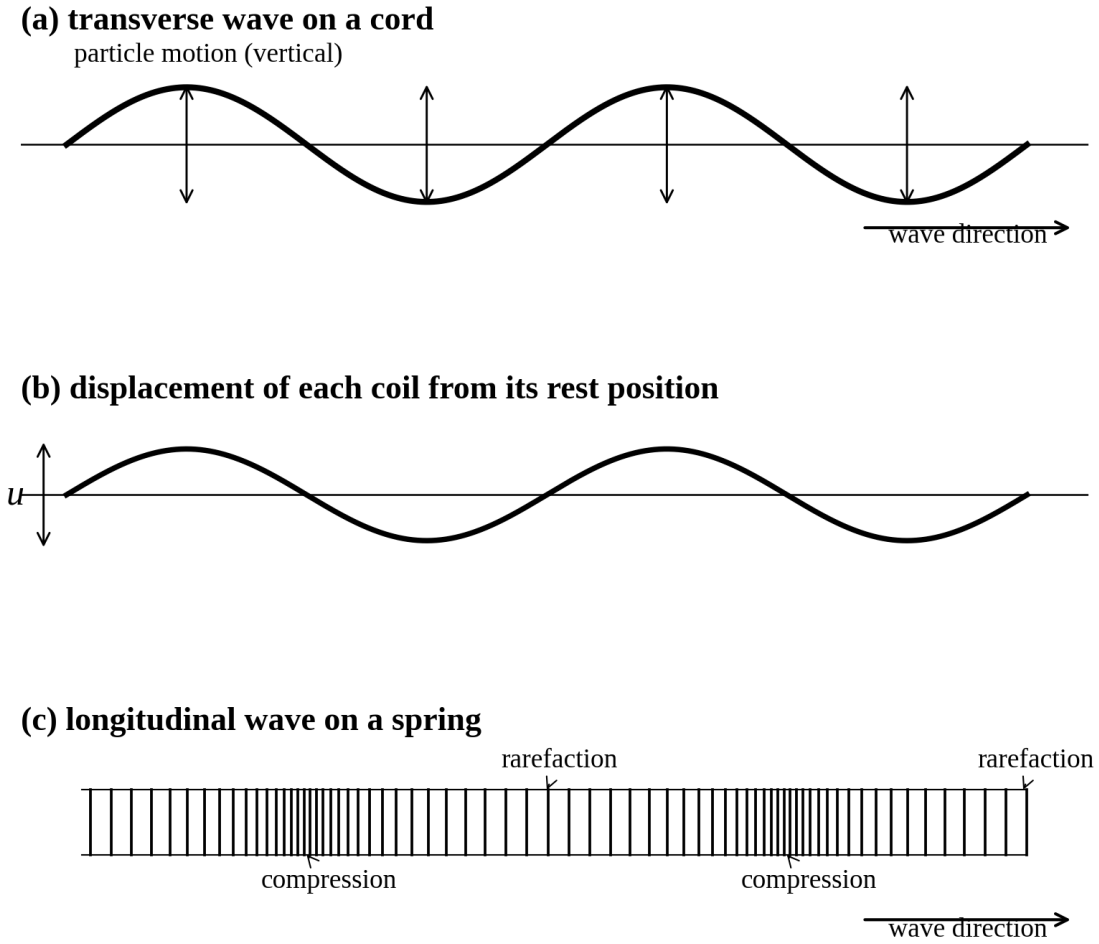


Figure: ***Transverse versus longitudinal.*** *(a) A transverse wave drawn on a stretched cord; particles move vertically while the wave moves horizontally. (b) The displacement of each coil of a spring from its rest position, plotted against position along the spring. (c) The same longitudinal wave drawn as the spring itself, with compressions where the displacement-curve slope is steeply negative and rarefactions where it is steeply positive.*

The wave vocabulary

Almost every question in this chapter, in any examination paper you will ever sit, will ask you to read or to compute one of the following six numbers. They are the wave's coordinates, in the same sense that latitude and longitude are a city's coordinates. If you are fluent in them, the chapter will feel light. If they are slippery, the chapter will be exhausting. Two minutes of careful attention now buys an enormous amount of comfort later.

Amplitude. The amplitude, written **A**, *is the maximum displacement of any particle of the medium from its equilibrium position during the wave's passage.* For a wave on water, A is the height of a crest above the still surface. For a wave on a guitar string, A is half the distance between the highest and lowest points of the vibrating string. The amplitude is always a positive number, and it has the units of length: metres or centimetres or, for a sound wave, micrometres.

Wavelength. The wavelength, written **λ** (Greek lambda), *is the distance between any two consecutive points in identical phase.* It may be: between two consecutive crests of a wave, between two consecutive troughs, or between any two consecutive points of the medium that are doing the same thing at the same time. For a **longitudinal wave**, the **wavelength** *is the distance between two consecutive compressions*. The wavelength is measured in metres.

Period. The period, written **T**, *is the time taken for any one particle of the medium to complete a single full oscillation:* up, down, and back to the start. The period is also *the time the wave pattern takes to advance by exactly* ***one wavelength.*** The period is measured in seconds.

Frequency. The frequency, written **f**, *is the number of full oscillations completed by a particle per second.* It is the reciprocal of the period:

$$\mathbf{f} = \frac{\mathbf{1}}{\mathbf{T}}$$

The frequency is measured in hertz, where 1Hz means one oscillation per second. A piano A above middle C has a frequency of 440Hz. The hum of a Tanzanian power line is at 50Hz. The cosmic microwave background, which is light left over from

the early universe, peaks at about 160GHz. Frequencies span an enormous range; periods are their inverses and span the inverse range.

Angular frequency. The angular frequency, written **ω** (Greek omega), *is the frequency expressed in radians per second* rather than oscillations per second. Since one full oscillation is 2π radians,

$$\omega = 2\pi f = \frac{2\pi}{T}$$

The angular frequency is measured in rad/s. *Why bother with two different ways of saying "frequency"?* Because the equation of a wave, like the equation of SHM in Chapter 6, contains a sine. The argument of a sine is an angle. Angles are measured in radians. So the natural variable for the time part of a wave's equation is ω, not f. The two are related by a factor of 2π and you must be fluent in both.

Wave number. The wave number, written **k**, is the spatial counterpart of the angular frequency. Where ω tells you how many radians of oscillation pass per second, k tells you *how many radians of oscillation fit into a metre*. Since one full wavelength is 2π radians of phase,

$$k = \frac{2\pi}{\lambda}$$

The wave number is measured in radians per metre, or simply in m^{-1}. A wave with a long wavelength (for example, an ocean swell, $\lambda \approx 100\text{m}$) has a small wave number ($k \approx 0.063\text{rad/m}$). A wave with a short wavelength (a fruit-bat echolocation call along the shores of Lake Tanganyika, $\lambda \approx 1\text{cm}$) has a large wave number ($k \approx 628\text{rad/m}$). The wave number is the spatial pulse-rate of a wave.

Wave speed. The sixth number is **v**, the wave speed: *the speed at which the pattern of disturbance travels through the medium.* The wave speed is **not** the speed of any particle. The wave speed is determined by the medium, the same way the speed of a transverse vibration on a string is determined by the tension and the linear mass density. From the geometry of the wave, the wave speed is related to the wavelength and the period by:

$$v = \frac{\lambda}{T} = f\lambda = \frac{\omega}{k}$$

All three forms are useful. The first (λ/T) is the most physical: in one period, the wave moves forward by one wavelength, so its speed is wavelength per period. The second (fλ) is the same thing reorganised so the experimentalist's two natural measurements (frequency and wavelength) appear directly. The third (ω/k) is the form in which v emerges naturally from the wave equation we will derive in the next section. They are the same equation in three coats.

Wave behaviour at a boundary: frequency is preserved, wavelength is not

One physical fact about waves is so important, and so frequently misunderstood by students, that we will state it now and refer back to it in three later sections. When a wave crosses from one medium into another (a sound wave going from air into water, an ultrasound pulse moving from fat into muscle, a seismic wave moving from sedimentary rock into harder basement rock), three things in general can happen to it. Some of the wave reflects off the boundary. Some of the wave transmits through. And the transmitted part travels at a different speed in the new medium, because the wave speed depends on the medium and the new medium has different elastic and inertial properties.

What does **not** change at a boundary is the wave's frequency. The reason is physical, not mathematical. The source on one side of the boundary is producing the wave at a particular rate, say 440 oscillations per second. Each oscillation must, for continuity, push exactly one corresponding oscillation across the boundary into the new medium. If the new medium produced more oscillations per second than the old medium delivered, oscillations would be appearing from nowhere; if it produced fewer, oscillations would be disappearing without trace. Neither is allowed by the conservation of disturbance count. The rate at which oscillations cross the boundary must be the same on both sides. Therefore the frequency is the same on both sides.

The wave speed, however, **does** change. And since $v = f\lambda$ with f kept constant, the wavelength has no choice but to scale in proportion to the speed. If the wave speed doubles in the new medium, the wavelength doubles. If it halves, the wavelength halves. This rule will resurface in the very last section of the chapter, when an ultrasound transducer pushes a 5MHz wave from fat (where it travels at about 1450m/s) into muscle (where it travels at about 1580m/s); we will use frequency invariance and $v = f\lambda$ to work out how the wavelength changes across the boundary, and the radiologist's image-resolution argument will follow from that. For now, simply remember the rule. If somebody asks what changes when a wave crosses from one medium into another, the truthful answer is: *the speed changes; the wavelength changes in step; the frequency does not.*

Two worked examples now. They are kinder than they look, but only the second one is keeping a straight face about it.

BINDER Example 1

On a calm morning at the southern end of Lake Tanganyika, a fisherman watches the small wooden float on his line bob steadily up and down. The crests of the surface waves on the lake are 2.4m apart, and the float completes 3 oscillations every 2 seconds. Find the wavelength, frequency, period, angular frequency, wave number, and wave speed of the surface waves.

Solution

The wavelength is given to us directly: it is the distance between consecutive crests.

$$\lambda = 2.4\text{m}$$

The frequency is the number of complete oscillations per second. The float completes 3 in 2 seconds, so:

$$\text{f} = \frac{3}{2\text{s}} = 1.5\text{Hz}$$

The period is the time for one full oscillation, which is the reciprocal of the frequency:

$$\text{T} = \frac{1}{\text{f}} = \frac{1}{1.5\text{Hz}} = 0.667\text{s}$$

The angular frequency is the frequency in radians per second:

$$\omega = 2\pi\text{f} = 2\pi \times 1.5\text{Hz} = 3\pi\,\text{rad/s} = 9.42\text{rad/s}$$

The wave number is 2π divided by the wavelength:

$$\text{k} = \frac{2\pi}{\lambda} = \frac{2\pi}{2.4\text{m}} = 2.62\text{rad/m}$$

And the wave speed comes from $v = f\lambda$:

$$\text{v} = \text{f}\lambda = 1.5\text{Hz} \times 2.4\text{m} = 3.6\text{m/s}$$

Making sense of the answer: *A surface-water wave on a calm lake travelling at 3.6m/s, or about 13km/h, is exactly what a casual eye would call "a gentle ripple cantering toward the shore". The number is right. So is its bookkeeping cousin: 1.5 oscillations per second is what a fisherman watching his float would describe as "bobbing comfortably". Six numbers, all of them describing the same wave, all of them computed from two pieces of given data: the crest spacing and the count of bobs in 2 seconds. That is the entire chapter's calculation pattern in miniature: read off the easy numbers, use $v = f\lambda$ or $\omega = 2\pi f$ or $k = 2\pi/\lambda$ to fill in the rest, and watch the units cancel cleanly.*

Think like a physicist: *Notice what just happened. The fisherman gave us only two numbers: the spacing between consecutive crests, and the rate at which his float bobbed. From those two numbers we extracted six. This is not a free lunch; it is the deep structure of a wave. A monochromatic wave has only two truly independent pieces of information, one belonging to space (any one of λ or k) and one belonging to time (any one of T, f, or ω). Everything else, including the wave speed itself, is derived.*

BINDER Example 2

A sound wave of wavelength λ travels in a first medium with velocity v. It enters a second medium in which the velocity is 4v. Determine the wavelength of the sound wave in the second medium.

Solution

The source frequency is unchanged across the boundary; conservation of oscillation count requires the same number of oscillations per second on both sides. Call that frequency f. In medium 1,

$$\text{v} = \text{f}\lambda \implies \text{f} = \frac{\text{v}}{\lambda}$$

In medium 2, the wave speed is 4v and the same frequency f applies. Calling the new wavelength λ', the basic relation gives:

$$4\text{v} = \text{f}\lambda' = \left(\frac{\text{v}}{\lambda}\right)\lambda' \implies \lambda' = 4\lambda$$

So the wavelength in the second medium is four times the wavelength in the first.

Making sense of the answer: *Notice that the answer is decided not by clever algebra but by a single physical commitment: the source dictates the rate; the medium dictates the speed. Wavelength is a derived quantity, $\lambda = v/f$, and it does whatever the speed and frequency tell it to do.*

Think like a physicist: *There is a beautiful continuity argument hidden inside this calculation. Imagine a long row of disturbance peaks marching toward the boundary in medium 1. They arrive at the boundary at a steady rate, f peaks per second. On the far side, they must depart at exactly the same rate, f peaks per second, for the same reason that cars on a road cannot vanish at a checkpoint. But on the far side they are spreading out four times as fast. Therefore on the far side the spacing between consecutive peaks must be four times as large: $\lambda' = 4\lambda$. The wavelength is just the geometric record of how far one peak has gone before the next one is born; if the peaks are now travelling four times as quickly, each peak runs four times as far before its successor turns up. This is the calculation done with hands instead of algebra, and the two views must agree. They do.*

With those two examples we have caught up, in vocabulary if not yet in equations, with the language the rest of the chapter will speak. We know what a wave is. We know its six descriptive numbers. We know that two of those numbers, one belonging to space and one belonging to time, are enough to pin down the rest. We know what happens when a wave crosses a boundary. What we do not yet have is the actual equation of a wave: the one Mr. Akilikubwa wrote in the wet sand and the lake then erased. The next section finally writes that equation, derives it from things you already know about simple harmonic motion, and teaches you to read its parameters off the page the way a musician reads notes off a stave. The chapter has, until now, been about words. From here on it will be about a single sentence in those words.

THE PROGRESSIVE WAVE EQUATION

So far we have spoken about a wave the way an ornithologist speaks about a bird that has flown overhead: in admiring generalities. We know what a wave does, what it carries, and what it does not carry. We know its six descriptive numbers. What we do not yet have is the equation that puts all of this on paper. The point of this section is to write that equation, derive it from things you already know, and then teach you to read it the way a musician reads a score.

The equation we are after is the same one Mr. Akilikubwa wrote in the wet sand at the start of the chapter. It will turn out to be a sentence in three pieces of vocabulary, all of which you already speak fluently: the amplitude A from this section's predecessor, the angular frequency ω from Chapter 6, and the wave number k from a few pages ago. The derivation rests on one physical observation: a wave is simple harmonic motion shared, with a delay, between every particle of the medium. Once we have the equation, every wave problem in the chapter becomes one of two things: read off its parameters, or build it from given parameters. There is no third option, and there is no third equation.

From SHM to a travelling wave: a derivation in two steps

Picture a long horizontal rope or a long stretched string, with one end held at the origin x = 0 by a person who can move that end up and down at will. Suppose she moves the end up and down in simple harmonic motion. From Chapter 6, we know that the displacement of that end, written y, is a sinusoidal function of time:

$$y(0,t) = A\sin(\omega t)$$

Here A is the amplitude of her hand, ω is its angular frequency, and the bracket $y(0,t)$ emphasises that this expression describes the displacement at the location x = 0 at any time t. So far this is just SHM with a label. Nothing new.

Now the new ingredient. The string is elastic. When the source end (at x = 0) moves up, it pulls its immediate neighbour up too, but only after a tiny delay. The neighbour pulls its neighbour, also after a delay. The motion runs along the string. By the time the disturbance reaches a point at distance x from the source, a finite time has elapsed. If the wave travels at speed v, that time is exactly the distance divided by the speed:

$$\text{delay} = \frac{x}{v}$$

So the particle at position x is doing the same simple harmonic motion as the source, but lagging it by a time of x divided by v. Whatever the source was doing at time t, the particle at x is doing at time t + x⁄v; equivalently, what the particle at x is doing right now (time t) is what the source was doing at the earlier time t − x⁄v. Substituting the earlier time into the SHM expression:

$$y(x,t) = A\sin\left(\omega\left(t - \frac{x}{v}\right)\right)$$

Distribute the angular frequency through the bracket:

$$y(x,t) = A\sin\left(\omega t - \frac{\omega}{v}x\right)$$

And the cluster ω⁄v has a name we have already met. The wave **number k** is, by definition, *the angular frequency divided by the wave speed:*

$$k = \frac{\omega}{v} \quad \text{so} \quad v = \frac{\omega}{k}$$

Substituting back, we have the equation that the chapter has been earning the right to write:

$$\mathbf{y(x,t) = A\sin(\omega t - kx)}$$

This is the displacement, at any time t and any position x, of a sinusoidal wave travelling along the string in the positive x-direction with amplitude A, angular frequency ω, and wave number k. The minus sign in front of kx is the wave's signature. It says: *the further along the string we look, the further behind the source the particle is in its cycle.* That phase lag, growing in step with x, is what makes the disturbance pattern shift to the right as time advances.

If the source is instead at the far end and the wave travels in the negative x-direction, the same derivation runs through with the opposite sign of the delay, and the equation reads:

$$\mathbf{y(x,t) = A\sin(\omega t + kx)}$$

So the sign of kx is the wave's compass*: a **minus** sign means a wave travelling toward **positive** x; a **plus** sign means a wave travelling toward **negative** x.* Read the sign first; everything else falls into place.

Reading the equation's parameters at a glance

The single most common type of question in this chapter, in any examination paper you will ever sit, will hand you a wave equation in the standard form and ask you to extract its parameters. The trick is to recognise that y = Asin(ωt − kx) is a recipe in which every parameter sits in plain sight. You read them off the way you read the time off a clock face, by knowing where each number lives.

Given y = Asin(ωt − kx), the parameters of the wave are:

amplitude = A (**the coefficient outside the sine**)

ω = **coefficient of t**,

k = **coefficient of x**

$$T = \frac{2\pi}{\omega}, \quad f = \frac{\omega}{2\pi}, \quad \lambda = \frac{2\pi}{k}, \quad v = \frac{\omega}{k} = f\lambda$$

Six parameters, five identities, one minute of work. There are no judgement calls and no algebraic tricks. Practise the extraction enough times that it becomes mechanical, because in an examination the marks will go to whoever does it without thinking.

Two velocities, one wave: keeping wave speed and particle speed apart

Before any worked examples, we must spend a careful minute on the conceptual error that this chapter is most ruthless about catching. The equation $y = A\sin(\omega t - kx)$ describes two motions at once.

The first is the motion of the wave pattern itself, which slides rigidly to the right at speed $v = \omega/k$.

The second is the motion of any individual particle of the string, which oscillates up and down without travelling anywhere on average.

These are different things. Their speeds are different numbers. The chapter that mistakes the two will fail.

To find the velocity of an individual particle at position x, we differentiate y with respect to t while holding x fixed. The particle's position depends on time, so its velocity is dy/dt:

$$\mathbf{v_p = \frac{dy}{dt} = A\omega\cos(\omega t - kx)}$$

This is also a sinusoid, but it is a sinusoid of a particle's velocity in time, not the wave's velocity. Its maximum value, achieved when the particle is passing through its equilibrium position, is:

$$\mathbf{v_{p,max} = A\omega}$$

Compare this to the wave speed. The wave speed $v = \omega/k$ is determined entirely by the medium (its tension, its linear mass density, its elasticity, and so on). It does not depend on the amplitude A. It does not depend on where on the string you look or on what the particle there happens to be doing. It is a property of the highway, not of any one car driving on it.

The maximum particle speed $v_{p,max} = A\omega$, on the other hand, depends on the amplitude. Make the wave taller and every particle on it moves faster up and down, even though the wave pattern itself slides past at exactly the same v as before. Doubling the amplitude doubles the maximum particle speed; it does not change the wave speed at all. The two are separate physical quantities, related to each other only by the equation that governs the wave, and asking for one when the question wants the other is the most reliable way to lose marks in this chapter. Examiners know this, and they place traps accordingly.

> **Warning: the two velocities of a wave**
>
> The **wave speed** $v = \omega/k$ is the speed at which the wave pattern moves along the medium. It is a property of the medium and does not depend on the amplitude.
>
> The **particle velocity** $v_p = dy/dt$ is the speed at which an individual particle is moving up and down at right angles to the medium. It depends on time, on position, and on the amplitude. Its maximum value is $v_{p,max} = A\omega$.
>
> These are completely different quantities. They have completely different values. Confusing them is the most common error in the chapter. In every numerical problem, ask first: *am I being asked for the speed of the pattern, or the speed of a particle?*

The differential equation that every progressive wave obeys

There is one further fact about $y = A\sin(\omega t - kx)$ that is worth establishing now, because it underpins half of the rest of the chapter and all of optics in Chapter 8. The displacement function we have written down is not just a static picture; it is a solution to a particular partial differential equation, called the wave equation:

$$\frac{\partial^2 y}{\partial t^2} = v^2 \frac{\partial^2 y}{\partial x^2}$$

Read this equation in plain English. On the left is how fast the displacement of a particle is accelerating in time. On the right is how sharply the wave profile is curving in space, scaled by the square of the wave speed. The equation says these two are equal, always. In a stretched string, this equation is just Newton's second law applied to a tiny segment of string, with restoring force from the tension on either side and inertia from the segment's linear mass density. We will not derive it from Newton here; we will simply check that $y = A\sin(\omega t - kx)$ is a solution.

Differentiate y twice with respect to t, holding x fixed:

$$\frac{\partial y}{\partial t} = A\omega\cos(\omega t - kx)$$

$$\frac{\partial^2 y}{\partial t^2} = -A\omega^2\sin(\omega t - kx) = -\omega^2 y$$

And twice with respect to x, holding t fixed:

$$\frac{\partial y}{\partial x} = -Ak\cos(\omega t - kx)$$

$$\frac{\partial^2 y}{\partial x^2} = -Ak^2\sin(\omega t - kx) = -k^2 y$$

Take the ratio of the two second derivatives:

$$\frac{\frac{\partial^2 y}{\partial t^2}}{\frac{\partial^2 y}{\partial x^2}} = \frac{-\omega^2 y}{-k^2 y} = \frac{\omega^2}{k^2} = \left(\frac{\omega}{k}\right)^2 = v^2$$

Hence:

$$\frac{\partial^2 y}{\partial t^2} = v^2 \frac{\partial^2 y}{\partial x^2}$$

So $y = A\sin(\omega t - kx)$ does indeed satisfy the wave equation, with the wave speed coming out as $v = \omega / k$. The arithmetic is honest. Two differentiations expose the deep structure: a wave is a function whose curvature in space is exactly proportional to its acceleration in time, with the proportionality constant being the square of the wave speed.

For now, give theory a short holiday and let us sharpen our understanding with six worked examples. None of them are unkind, but the third one quietly pretends to be a parameter-extraction problem and is in fact a particle-velocity problem in disguise. Read the question twice.

BINDER Example 3

A transverse wave is described by the equation $y(x,t) = 8\sin(200\pi t - 4\pi x)$ mm, where t is in seconds and x is in metres. Find the amplitude, angular frequency, wave number, period, frequency, wavelength, and wave speed.

Solution

Compare the equation directly with the standard form $y = A\sin(\omega t - kx)$.

$$A = 8.0\text{mm} = 8.0 \times 10^{-3}\text{m}$$

$$\omega = 200\pi\ \text{rad/s}$$

$$k = 4\pi\ \text{rad/m}$$

Now compute the period, frequency, wavelength, and wave speed from the standard relations.

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{200\pi\ \text{rad/s}} = 0.010\text{s}$$

$$f = \frac{1}{T} = \frac{1}{0.010\text{s}} = 100\text{Hz}$$

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{4\pi\ \text{rad/m}} = 0.50\text{m}$$

$$v = \frac{\omega}{k} = \frac{200\pi\ \text{rad/s}}{4\pi\ \text{rad/m}} = 50\text{m/s}$$

Making sense of the answer: *A 50m/s transverse wave on a stretched wire is well within the range of physical realism, somewhere between a guitar string at modest tension and a clothesline carrying a hand pulse. A 100Hz oscillation is a low musical note, well below the lower edge of human hearing if it were a sound. An 8mm amplitude is a generously plucked string; you would see it. Notice the bookkeeping: every one of the six numbers came from comparing the given equation against $y = A\sin(\omega t - kx)$ and applying $T = 2\pi/\omega$, $f = 1/T$, $\lambda = 2\pi/k$, $v = \omega/k$. There were no judgement calls. This is the engine of every parameter-extraction problem in the chapter.*

Think like a physicist: *Look once more at the wave speed: $v = \omega/k$. We could equally well have computed it from $v = f\lambda = (100\text{Hz})(0.50\text{m}) = 50\text{m/s}$. Both routes give the same answer because the four numbers ω, k, f, λ are not independent: $\omega = 2\pi f$ and $k = 2\pi/\lambda$ pin them together. There are only two genuinely independent pieces of information in this wave (one belonging to time, one belonging to space), and once you have either pair (T, λ), (f, λ), or (ω, k), every other parameter follows. This is the lesson Example 1 hinted at in the previous section, now made unambiguous: a monochromatic wave has two degrees of freedom, not six.*

BINDER Example 4

A horizontal stretched elastic string of length 3.0m and mass 12kg is subjected to a tension of 1.6N. A transverse wave of frequency 40Hz is propagated down the string. Determine the distance between successive crests of this wave.

Solution

The distance between successive crests is the wavelength. We need v first, then λ from $v = f\lambda$. The linear mass density is

$$\mu = \frac{m}{L} = \frac{12\text{kg}}{3.0\text{m}} = 4.0\text{kg/m}$$

The wave speed for a transverse wave on a stretched string was derived at the end of Chapter 6 as:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{1.6\text{N}}{4.0\text{kg/m}}} = 0.632\text{m/s}$$

The wavelength then follows from $v = f\lambda$:

$$\lambda = \frac{v}{f} = \frac{0.632\text{m/s}}{40\text{Hz}} = 0.0158\text{m} \approx 1.6\text{cm}$$

Making sense of the answer: *A heavy string under a tiny tension carries waves slowly, and a slow wave at a moderate frequency has a short wavelength. The numbers obey the right physics: μ = 4kg/m is unusually large (most laboratory strings are a thousand times lighter per metre), T = 1.6N is unusually small (most laboratory strings are tensioned to tens or hundreds of newtons), so a wave speed of 0.6m/s is exactly the order of magnitude expected. The 1.6cm wavelength is what you would see if you happened to film the string in slow motion. The arithmetic is unforgiving, but it is also unambiguous: get μ right, get v right, get λ from v = fλ.*

Think like a physicist: *This problem hands you four numbers (L, m, T, f) and asks you for one (λ). It looks unbalanced, but it is not. The four given numbers split into two functional pairs. The pair (L, m) gives μ. The pair (T, μ) gives v. Then (v, f) gives λ. Each pair contributes one of the equation links of the chapter: μ = m/L, v = √(T/μ), v = fλ. Three links, three steps, one answer. Examination problems often look as if they are giving you more information than you need; they are not. Every given number has a job. If a number seems irrelevant, look for a chain of equations you have not noticed yet.*

BINDER Example 5

Write the equation of a progressive wave travelling in the positive x-direction with amplitude 0.030m, frequency 80Hz, and wave speed 40m/s.

Solution

This is the reverse of Example 3: from given physical parameters, build the wave equation. The standard form is $y = A\sin(\omega t - kx)$. We need ω and k.

$$\omega = 2\pi f = 2\pi \times 80\text{Hz} = 160\pi\ \text{rad/s}$$

$$\lambda = \frac{v}{f} = \frac{40\text{m/s}}{80\text{Hz}} = 0.50\text{m}$$

$$k = \frac{2\pi}{\lambda} = \frac{2\pi}{0.50\text{m}} = 4\pi\ \text{rad/m}$$

Substituting A = 0.030m, ω = 160π rad/s, and k = 4π rad/m into $y = A\sin(\omega t - kx)$:

$$\mathbf{y(x,t) = 0.030\sin(160\pi t - 4\pi x)\ m},$$

with t in seconds and x in metres.

Making sense of the answer: *Read the equation back to confirm. Coefficient of t is 160π rad/s; dividing by 2π gives 80Hz. Coefficient of x is 4π rad/m; 2π divided by it gives 0.50m. Wave speed is the ratio 160π/4π = 40m/s. All three given numbers come back out unchanged. The amplitude was carried straight across. The minus sign in front of kx encodes "travelling in the positive x-direction", which is what the question asked for. Every check passes.*

Think like a physicist: *Construction problems like this one are how you debug your own understanding. If you can build a wave equation from physical parameters, you understand what each symbol means; if you cannot, the symbols are still words you have not internalised. Keep practising the round trip in both directions. Extraction problems like Example 3 take you from the equation to the physics; construction problems like this one take you from the physics back to the equation. They are the same skill viewed from opposite sides, and an honest physicist owns both.*

BINDER Example 6

A transverse wave on a string is given by $y(x,t) = 0.040\sin(10\pi t - 2\pi x)$ m, where t is in seconds and x is in metres. (a) Find the speed of the wave. (b) Find the velocity of the particle at x = 0.125m at time t = 0.050s. (c) Find the maximum particle speed anywhere on the string. (d) Explain physically why the answers to (a) and (b) are different.

Solution

Comparing with $y = A\sin(\omega t - kx)$ we read off A = 0.040m, ω = 10π rad/s, k = 2π rad/m.

(a) The wave speed is the ratio ω/k:

$$v = \frac{\omega}{k} = \frac{10\pi\ \text{rad/s}}{2\pi\ \text{rad/m}} = 5.0\text{m/s}$$

(b) The particle velocity is the time derivative of y at fixed x:

$$v_p = \frac{\partial y}{\partial t} = A\omega\cos(\omega t - kx)$$

First evaluate the cosine argument at the chosen point. With x = 0.125m and t = 0.050s,

$$\omega t - kx = (10\pi\ \text{rad/s})(0.050\text{s}) - (2\pi\ \text{rad/m})(0.125\text{m}) = 0.5\pi - 0.25\pi = \frac{\pi}{4}\ \text{rad}$$

The cosine of π/4 is √2/2. Substituting:

$$v_p(0.125\text{m}, 0.050\text{s}) = A\omega\cos\left(\frac{\pi}{4}\right) = (0.040\text{m})(10\pi\ \text{rad/s})\left(\frac{\sqrt{2}}{2}\right) = 0.20\pi\sqrt{2}\ \text{m/s} \approx 0.889\text{m/s}$$

(c) The particle speed is maximised when the cosine is ±1, giving:

$$v_{p,max} = A\omega = (0.040\text{m})(10\pi\ \text{rad/s}) = 0.40\pi\ \text{m/s} \approx 1.26\text{m/s}$$

So at the chosen point and instant the particle is moving up at about 0.889m/s, which is $\sqrt{2}/2$ (about 71%) of its maximum particle speed of 1.26m/s. Different particles, at other positions and instants, will be moving at different speeds between zero and the maximum, depending only on the cosine of their particular phase $\omega t - kx$.

(d) The wave speed and the particle speed are different physical quantities. The wave speed v = 5.0m/s is the speed at which the disturbance pattern moves to the right along the string; it is determined by the medium. The particle velocity $v_p \approx 0.889$m/s is the speed at which the particle at (x = 0.125m, t = 0.050s) is moving up at that instant; it is determined by the equation that governs the wave evaluated at that one point and one time.

Understand: The two numbers describe different physical things: one is a property of the highway (the medium), the other is a property of one car driving on the highway at one moment (one particle at one phase).

Making sense of the answer: *Notice the order of magnitude. The wave pattern is moving along the string at 5.0m/s, while the particle at our chosen point is wiggling upward at about 0.89m/s, somewhat below its own maximum of 1.26m/s. The disturbance is racing past the particle. This is the right qualitative picture: a wave is a fast-moving pattern hosted by particles each of which is doing comparatively modest oscillation in place. Different particles, at other positions or instants, will be moving at different speeds, between zero and the maximum, but never at the wave speed itself. If the wave moved at the same speed as the particles, it would not be a wave at all; it would be a rigid translation of the medium.*

Think like a physicist: *There is an extreme case worth keeping at the back of your mind. The maximum particle speed $v_{p,max} = A\omega$ is proportional to the amplitude. If you crank the amplitude up high enough, you can in principle make the maximum particle speed exceed the wave speed $v = \omega/k$. When that happens, the small-amplitude approximation behind $y = A\sin(\omega t - kx)$ breaks down, and the wave starts to behave non-linearly: it can steepen, it can break, and it can shed energy through means we have not yet introduced. For ordinary waves at ordinary amplitudes we are far from this regime. The inequality $A\omega < v$ is the silent assumption that makes the linear theory work. For our wave, $A\omega = 1.26$m/s sits comfortably below $v = 5.0$m/s, so the linear theory holds.*

REAL Example 7

On the morning of 7 June, a seismograph operated by the Tanzania Meteorological Authority near Dodoma records the ground displacement of a small earthquake originating beneath the East African Rift. The recorded displacement is well-fit by the equation $y(x,t) = 0.0030\sin(3.0t - 0.0060x)$ m, where t is in seconds and x is in metres. From the equation alone, find (a) the amplitude of ground shaking, (b) the wave speed, (c) the wavelength, and (d) the frequency.

Solution

Compare with $y = A\sin(\omega t - kx)$. The amplitude appears as the coefficient outside the sine; the angular frequency is the coefficient of t; the wave number is the coefficient of x.

$$A = 0.0030\text{m} = 3.0\text{mm}$$
$$\omega = 3.0\ \text{rad/s}$$
$$k = 0.0060\ \text{rad/m}$$

(a) The amplitude is 3.0mm.

(b) The wave speed is the ratio:

$$v = \frac{\omega}{k} = \frac{3.0\ \text{rad/s}}{0.0060\ \text{rad/m}} = 500\text{m/s}$$

(c) The wavelength comes from $\lambda = 2\pi/k$:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{0.0060\ \text{rad/m}} \approx 1047\text{m} = 1.047\text{km}$$

(d) The frequency comes from $f = \omega/(2\pi)$:

$$f = \frac{\omega}{2\pi} = \frac{3.0\ \text{rad/s}}{2\pi} = 0.48\text{Hz}$$

Making sense of the answer: *A ground-shaking amplitude of 3mm is small, the kind of tremor that rattles cups and is felt by people sitting still but not by people walking; it is consistent with a small to moderate East African Rift event recorded at moderate epicentral distance. A wave speed of 500m/s is squarely in the range expected for seismic surface waves through Tanzanian sedimentary cover; this is roughly half the P-wave speed in solid rock, which is what we would expect for a shallow crustal disturbance. The wavelength of about a kilometre and the frequency of about half a hertz both fit the picture: long, slow undulations of the kind a seismograph records and a casual observer feels as a gentle swaying of buildings rather than a sharp jolt.*

Think like a physicist: *The three numbers (0.0030, 3.0, 0.0060) and the implicit minus sign are an extraordinary compression of physical information. From three numbers we extracted the depth of the disturbance (amplitude), the rate of the shaking (frequency), the spatial scale of the disturbance (wavelength), and the speed of its travel through the rock (wave speed). The seismograph trace is, in one literal sense, a sentence in the language of y = Asin(ωt − kx), and the language is rich enough that a single recorded equation tells the whole story. This is why physicists insist on writing waves as equations rather than as pictures. The picture is beautiful; the equation is computable.*

HOT Example 8

A transverse wave on a string is described by y(x,t) = 0.050 sin(100πt − 5πx) m, where t is in seconds and x is in metres. The string has linear mass density μ = 4.0 × 10⁻³ kg/m. (a) Find all six wave parameters (A, ω, k, T, f, λ) and the wave speed. (b) Find the tension in the string. (c) Find the velocity of the particle at x = 0.10m at time t = 0.010s. (d) Show that y satisfies the wave equation $\partial^2y/\partial t^2 = v^2 \cdot \partial^2y/\partial x^2$.

Solution

(a) Comparing with y = Asin(ωt − kx):

$$A = 0.050\text{m}, \quad \omega = 100\pi \text{ rad/s}, \quad k = 5\pi \text{ rad/m}$$

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{100\pi \text{ rad/s}} = 0.020\text{s}$$

$$f = \frac{1}{T} = 50\text{Hz}$$

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{5\pi \text{ rad/m}} = 0.40\text{m}$$

$$v = f\lambda = (50\text{Hz})(0.40\text{m}) = 20\text{m/s}$$

(b) The tension follows from $v = \sqrt{(T/\mu)}$ rearranged:

$$T = \mu v^2 = (4.0 \times 10^{-3}\text{ kg/m})(20\text{m/s}^2) = 1.6\text{N}$$

(c) The particle velocity is dy/dt at fixed x:

$$v_p = A\omega\cos(\omega t - kx)$$

At x = 0.10m, t = 0.010s, the argument of the cosine is

$$\omega t - kx = (100\pi \text{ rad/s})(0.010\text{s}) - (5\pi \text{ rad/m})(0.10\text{m}) = \pi - 0.5\pi = 0.5\pi \text{ rad}$$

$$\cos(0.5\pi) = 0$$

$$v_p = A\omega \times 0 = 0\text{m/s}$$

That particle, at that instant, is momentarily at rest. It is at one extreme of its oscillation, about to reverse direction.

(d) The wave equation we are checking is

$$\frac{\partial^2 y}{\partial t^2} = v^2 \cdot \frac{\partial^2 y}{\partial x^2}$$

We check this by differentiating the given equation y = 0.050 sin(100πt − 5πx) m directly, in three small stages: first the time side, then the space side, then the comparison.

The time side

Differentiate y once with respect to t, holding x fixed:

$$\frac{\partial y}{\partial t} = 0.050 \times 100\pi \times \cos(100\pi t - 5\pi x) = 5\pi\cos(100\pi t - 5\pi x)\text{ m/s}$$

Differentiate once more:

$$\frac{\partial^2 y}{\partial t^2} = -5\pi \times 100\pi \times \sin(100\pi t - 5\pi x) = -500\pi^2 \sin(100\pi t - 5\pi x)\text{ m/s}^2$$

Now notice that 0.050 sin(100πt − 5πx) is exactly y. We can therefore pull the factor 0.050 back inside the sine to expose y:

$$\frac{\partial^2 y}{\partial t^2} = -(100\pi)^2 \times y = -10000\pi^2\, y$$

The space side

Now differentiate y twice with respect to x, holding t fixed:

$$\frac{\partial y}{\partial x} = 0.050 \times (-5\pi) \times \cos(100\pi t - 5\pi x) = -0.25\pi\cos(100\pi t - 5\pi x)$$

$$\frac{\partial^2 y}{\partial x^2} = (-0.25\pi) \times (-5\pi) \times \big(-\sin(100\pi t - 5\pi x)\big) = -1.25\pi^2 \sin(100\pi t - 5\pi x)$$

Pull the 0.050 back inside the sine again to expose y:

$$\frac{\partial^2 y}{\partial x^2} = -(5\pi)^2 \times y = -25\pi^2\, y$$

The comparison

Substitute both results into the wave equation:

The left-hand side is;

$$\frac{\partial^2 y}{\partial t^2} = -10000\pi^2\, y$$

The right-hand side is;

$$v^2 \times \frac{\partial^2 y}{\partial x^2} = v^2 \times (-25\pi^2\, y) = -25\pi^2 v^2\, y$$

The two sides agree provided that

$$10000\pi^2 = 25\pi^2 v^2$$

Solve for v:

$$v^2 = \frac{10000}{25} = 400, \text{so } v = 20\text{m/s}.$$

This matches the wave speed v = 20m/s found in part (a). The wave equation is satisfied, and the wave speed that falls out of the differential equation is precisely the same wave speed we read directly from the ratio ω⁄k of the original equation. The arithmetic closes the loop.

Making sense of the answer: *All three parts of this problem fit together. The single given equation tells us everything about the wave (part a) and, combined with the linear mass density of the string, lets us deduce the tension (part b). The particle velocity (part c) zooms in on one particle at one instant and finds it momentarily at rest, exactly the point in its oscillation where its kinetic energy is least and its potential energy is greatest.*

Think like a physicist: *Look at the structure of part (d). Two derivatives in time gave a factor of $-(100\pi)^2 = -\omega^2$. Two derivatives in space gave a factor of $-(5\pi)^2 = -k^2$. The wave equation says these two are related by a factor of v^2. Squaring $v = \omega/k$ gives $v^2 = \omega^2/k^2$, and the ratio of the second derivatives reproduces this exactly: $10000\pi^2 \div 25\pi^2 = 400 = v^2$. This is more than a calculation; it is the reason the wave equation has the form it does. Acceleration in time and curvature in space are both proportional to y, with proportionality constants ω^2 and k^2, and their ratio fixes the wave speed. If you ever forget the wave equation, write down any sinusoidal wave you happen to know, differentiate twice each way, take the ratio, and the wave equation falls out of the algebra without your having to memorise it. This is what physicists mean when they say a structural equation is its own mnemonic.*

Eight worked examples, six wave parameters, and one master equation. We have done the heavy lifting: the equation y = Asin(ωt − kx) is no longer a mystery to be unpacked but a sentence whose vocabulary you can read at a glance. What we have not yet done is something that, on the surface, sounds simpler than anything in this section. We have not yet compared two waves. When two waves meet, what matters is not the value of either of their phases in isolation but the difference between them. The next section is about that difference.

PHASE, PHASE DIFFERENCE, AND WAVE INTENSITY

In the previous section we learned to read the equation of one wave. We can extract its parameters, build it from given physical quantities, and prove that it satisfies the wave equation. The next natural question, and the one this section answers, *is what to do when there are **two** of them.* Two waves on the same string. Two ripples meeting on Lake Victoria. Two voices arriving at the same ear from two different speakers. The numerical comparison of one wave to another rests almost entirely on a single concept: *the difference in phase between them.* Master phase difference and you have unlocked, in one sitting, the calculation that drives Young's double-slit experiment, the colours of an oil film on a puddle, the entire physics of antenna arrays, and (in less glamorous form) the noise-cancelling headphones a Form Five student might wear on the bus to school.

This section also picks up the second important property of a wave that the previous section postponed: the energy a wave carries. Every wave has an amplitude, and every amplitude implies a definite quantity of energy per unit volume of the medium. When that energy spreads outward from a point source through three-dimensional space, it dilutes in a strict geometrical pattern called the inverse-square law. We will derive that pattern, see what it implies about how the amplitude falls off with distance, and use it to read the loudness of any sound source from any distance.

Two ideas in this section, therefore: **phase difference** and **intensity**. They are linked by a common ancestor, the equation y = Asin(ωt − kx), which tells you both where in the cycle each particle is and how much energy each particle is carrying.

Phase: the angular bookkeeping of where you are in the cycle

Take the wave equation we earned in the previous section,

$$y = A\sin(\omega t - kx)$$

and look closely at the argument inside the sine, namely ωt − kx. This argument has a name: it is called the **phase** of the wave at position x at time t, often written φ. **Phase** *is an angle, measured in radians, that tells you exactly where in its cycle the particle at (x, t) is at that moment.*

- A phase of 0 means the particle is at its equilibrium position and moving upward.
- A phase of π⁄2 means the particle is at its highest point.
- A phase of π means it is back at equilibrium, but now moving downward.
- A phase of 3π⁄2 means it is at its lowest point.

- After a phase advance of 2π the particle returns to where the cycle started and the whole story repeats.

So a wave's equation is really a bookkeeping system: every particle of the medium has its own phase clock, and the equation $y = A\sin(\omega t - kx)$ is the master rule that tells you what each clock reads at each instant.

Two observations follow immediately.

1) The phase of a single particle (held at a fixed x) advances steadily in time at a rate ω. Time t increases, kx is constant, so $\varphi = \omega t - kx$ grows linearly.
2) At any single instant in time, two particles at different positions read different phases, because their kx values are different. This second observation is the entire content of phase difference, which is what the next subsection is about.

Phase difference between two points on the same wave

Pick two specific points on the same wave at the same instant t. Call them point 1 (at position x_1) and point 2 (at position x_2). Each has its own phase at that instant:

$$\varphi_1 = \omega t - kx_1, \quad \varphi_2 = \omega t - kx_2$$

The **phase difference** between point 2 and point 1 is, by definition, the difference of their phases:

$$\Delta\varphi = \varphi_1 - \varphi_2 = k(x_2 - x_1) = k\Delta x \text{ (where } \Delta x \text{ is the distance separating the two points)}$$

The ωt cancels because it is shared by both points: a phase difference is taken at the same instant. The remaining piece, kΔx, is the entire result. Substituting $k = 2\pi/\lambda$:

$$\boldsymbol{\Delta\varphi = \frac{2\pi}{\lambda}\Delta x}$$

This is the master equation of this section**. In words:** *the phase difference between two points on a wave is 2π radians per wavelength of separation.* Two points exactly one wavelength apart have a phase difference of 2π (which means they are doing the same thing at the same time, just one cycle apart). Two points exactly half a wavelength apart have a phase difference of π (which means one is doing the exact opposite of the other, called **antiphase**). Two points exactly a quarter wavelength apart have a phase difference of π/2 (which means one is at a crest while the other is crossing equilibrium).

Three landmark cases worth committing to memory:

- $\Delta\varphi = 0$ (in phase, doing the same thing at the same instant)
- $\Delta\varphi = \pi$ (in antiphase, doing the exact opposite)
- $\Delta\varphi = 2\pi$ (in phase again, separated by one full wavelength)

Any phase difference between these landmarks describes a state somewhere between full agreement and full disagreement. The mathematics treats them all on the same footing, but physically the in-phase and antiphase cases are the most important, because they are the cases in which the next section's discussion of superposition produces the cleanest results.

Path difference and phase difference: two faces of the same idea

The quantity Δx in the master equation has a special name of its own. It is called the **path difference** between the two points. Path difference is measured in metres; phase difference is measured in radians. They are two ways of describing the same physical separation, related by the bookkeeping rule:

$$\textbf{phase difference} = \frac{2\pi}{\lambda} \times \textbf{path difference}$$

Why introduce the same idea twice with two different names? Because in some contexts the natural way to describe a separation is in physical units of length (a path difference of 0.85m), and in other contexts the natural way is in angular units of phase (a phase difference of π radians). The two are interchangeable, and we will move freely between them. In Chapter 8, when we calculate the bright and dark fringes in Young's double-slit experiment, the natural quantity will be a path difference of one wavelength or half a wavelength. The same equation we just wrote will turn each path difference into the corresponding phase difference, and the phase difference will then determine whether two light waves arriving at a given point on a screen reinforce each other or cancel each other out.

Hold on to the equation. It is small, and it is everywhere.

The energy a wave carries: intensity and the inverse-square law

We have spent this chapter so far on the kinematics of a wave: where each particle is, how fast it is moving, what its phase is. We now turn to dynamics. A wave carries energy. The amount of energy depends on the amplitude in a particular way, and when that energy spreads through space from a point source it dilutes in a particular geometric pattern. Two equations cover this whole subsection.

We start with the energy. From Chapter 6, the total mechanical energy of any simple harmonic oscillator is

$$E = \frac{1}{2}m\omega^2A^2$$

Every particle on a wave is doing simple harmonic motion at angular frequency ω with amplitude A, so every particle carries an energy proportional to A^2. When we accumulate the energies of all the particles in a unit volume of the medium, the energy

density of the wave (energy per cubic metre) is also proportional to A^2. And when we ask how much of this energy crosses a unit area per unit time, the same proportionality survives: the **intensity** of a wave, written I, is proportional to A^2:

$$I \propto A^2$$

Intensity has the SI units watts per square metre (W/m^2), the same units as solar flux. It is the density of power flow through space carried by the wave, and the proportionality with A^2 has remarkable physical consequences. Doubling the amplitude of a wave quadruples its intensity. Halving the amplitude quarters the intensity. The square law is unforgiving and is the reason a softly-strummed guitar is, in energy terms, dramatically less loud than a vigorously strummed one, even though the amplitude looks only modestly larger to the eye.

Now consider a **point source** of waves: a small object that radiates energy outward equally in all directions. A small loudspeaker, a small lamp, a small earthquake epicentre. By symmetry the wave fronts emanating from such a source are concentric spheres. If the source emits total power P (in watts), and the medium does not absorb energy as the wave passes through it, then the entire power P must cross every spherical surface centred on the source, however far away. The surface area of a sphere of radius r is $4\pi r^2$, so the intensity at distance r is

$$I = \frac{P}{4\pi r^2}$$

This relationship is called the **inverse-square law** of intensity. The **intensity at distance r** *is the source's total power divided by the surface area of a sphere of radius r.* Doubling the distance multiplies the area by 4 and divides the intensity by 4. Tripling the distance multiplies the area by 9 and divides the intensity by 9. Move twice as far from a uniform sound source and the loudness, in energy terms, drops to a quarter of what it was.

Now combine the two ideas. Intensity is proportional to A^2, and intensity falls as $1/r^2$. So A^2 falls as $1/r^2$, which means amplitude itself falls as $1/r$. Doubling the distance halves the amplitude (not quarters it!) but quarters the intensity. The two distance dependencies are different, and the difference is a square root. This is the trap students fall into most often. Memorise the two rules side by side:

$$\mathbf{I \propto \frac{1}{r^2}} \quad \text{and} \quad \mathbf{A \propto \frac{1}{r}}$$

Or in ratio form, between two different distances r_1 and r_2 from the same source:

$$\mathbf{\frac{I_1}{I_2} = \left(\frac{r_2}{r_1}\right)^2} \quad \text{and} \quad \mathbf{\frac{A_1}{A_2} = \frac{r_2}{r_1}}$$

The ratio of intensities is the square of the inverse ratio of distances; the ratio of amplitudes is just the inverse ratio of distances. The third worked example below is built precisely on this distinction.

Three worked examples now. Pen ready

BINDER Example 9

A progressive wave is described by the equation $y = 0.04\sin(200\pi t - (20/17)\pi x)$ m, where t is in seconds and x is in metres. Find the phase difference between (a) the points at x = 0.25m and x = 1.10m, and (b) the points at x = 0 and $x = \lambda/4$.

Solution

Comparing the equation with the standard form $y = A\sin(\omega t - kx)$, we read off:

$$\omega = 200\pi \text{ rad/s}, \quad k = \frac{20}{17}\pi \text{ rad/m}$$

(a) The path difference between the two given points is

$$\Delta x = x_2 - x_1 = 1.10\text{m} - 0.25\text{m} = 0.85\text{m}$$

Multiply by k to get the phase difference:

$$\Delta\varphi = k\Delta x = \left(\frac{20}{17}\pi \text{ rad/m}\right) \times (0.85\text{m}) = \pi \text{ rad}$$

(b) When the path difference is exactly $\lambda/4$, the phase difference is

$$\Delta\varphi = \frac{2\pi}{\lambda} \times \frac{\lambda}{4} = \frac{\pi}{2} \text{ rad}$$

The wavelength λ cancels, so the answer does not depend on what λ actually is. Two points separated by a quarter of a wavelength are always $\pi/2$ radians out of phase, on any wave whatsoever.

Making sense of the answer: *In part (a), the two specific points x = 0.25m and x = 1.10m are exactly half a wavelength apart on this particular wave. Verify it: the wavelength is $\lambda = 2\pi/k = 2\pi \div ((20/17)\pi) = 17/10 = 1.7$m, so half a wavelength is 0.85m, matching Δx exactly. Two points half a wavelength apart are in antiphase, which is what $\Delta\varphi = \pi$ says. The cleverness of the question is that the numbers 0.25, 1.10, and 20/17 were chosen so that the fractions cancel without leaving a numerical residue. Examiners do this on purpose. Trust the cancellation; do not try to convert to decimals halfway through. In part (b), the answer falls out without computing λ at all, because the same wavelength appears in numerator and denominator. The lesson: when a path difference is given as a fraction of λ, the corresponding phase difference is the same fraction of 2π. No arithmetic required.*

Think like a physicist: *Phase difference is a dimensionless angle, but path difference is a length. The conversion between them, $\Delta\varphi = (2\pi/\lambda)\Delta x$, is the only equation in this section that you must memorise. Notice its structure: the right-hand side multiplies a length (Δx) by an inverse length ($2\pi/\lambda$ has units of rad/m), and the result is an angle (in radians). Units are doing real work here. If you ever forget the formula, restore it by asking:* ***how many radians of phase fit into one wavelength?*** *Answer: 2π. So per metre of separation, the phase advances by $2\pi/\lambda$. Multiply by the actual separation Δx to get the actual phase difference. You have just rederived the formula on the spot, and that is the kind of confidence that survives an examination room.*

REAL Example 10

A small wooden buoy on Lake Victoria, used by a Mwanza fisherman to mark the position of his fishing nets, rises and falls in simple harmonic motion with a period of 2.0s as the lake's surface waves pass through it. A second buoy of identical design floats 5.0m away in the direction the waves are travelling. The fisherman notices that the second buoy reaches its highest point exactly 0.5s after the first buoy reaches its own highest point. From this single observation, find (a) the phase difference between the oscillations of the two buoys, (b) the wave speed and the wavelength of the surface waves, and (c) verify the phase difference from part (a) using the master equation $\Delta\varphi = (2\pi/\lambda)\Delta x$ with the wavelength from part (b).

Solution

(a) Both buoys oscillate at the same period T = 2.0s, but the second buoy lags the first by a time $\Delta t = 0.5s$. A particle's phase advances by 2π in one full period, so in a time Δt it advances by the fraction $\Delta t/T$ of 2π. Therefore the second buoy lags the first in phase by:

$$\Delta\varphi = \frac{\Delta t}{T} \times 2\pi = \frac{0.5s}{2.0s} \times 2\pi = \frac{\pi}{2}\ \text{rad}$$

(b) The wave is moving at a speed such that whatever it does at the first buoy is repeated 0.5s later at the second buoy, 5.0m downstream. So the wave covers 5.0m in 0.5s:

$$v = \frac{\Delta x}{\Delta t} = \frac{5.0m}{0.5s} = 10m/s$$

The wavelength then follows from $\lambda = vT$:

$$\lambda = vT = (10m/s)(2.0s) = 20m$$

(c) Substitute the path difference $\Delta x = 5.0m$ and the wavelength $\lambda = 20m$ into the master equation:

$$\Delta\varphi = \frac{2\pi}{\lambda}\Delta x = \frac{2\pi}{20m} \times 5.0m = \frac{\pi}{2}\ \text{rad}$$

Same answer. The two routes: one through time delay and period, the other through path difference and wavelength; agree exactly, as they must, because they are describing the same physical lag in two different units.

Making sense of the answer: *The relationship between time delay and phase difference is a natural one. A particle's phase advances by 2π in one period T, so in any time Δt it advances by $(\Delta t/T)\cdot 2\pi$. Half a period of delay corresponds to π radians (antiphase). A quarter period corresponds to $\pi/2$ (which is what we found here). A full period corresponds to 2π (back in phase). The two buoys 5m apart on Lake Victoria are reading the same wave at different stages of the same cycle, and 5m happens to be exactly a quarter of the 20m wavelength. The numbers in this problem are chosen so the phase difference works out clean; in any real measurement the fisherman would more likely find a non-textbook fraction.*

Think like a physicist: *Notice what the fisherman did* ***not*** *need to do. He did not measure the wavelength directly (no ruler stretched over the lake's surface). He did not measure the wave speed directly (no tracked crest with a stopwatch). He measured a time delay (0.5s), a separation (5.0m), and his own buoy's bobbing period (2.0s). From those three numbers and the master equation $\Delta\varphi = (2\pi/\lambda)\Delta x$, plus the relation $\lambda = vT$, he extracted both the phase difference and every other parameter of the wave. This is the same two-degrees-of-freedom point we made earlier: a wave is fixed by two pieces of information, and once you have any two (here, period and wave speed deduced from the time-and-distance pair), every other parameter follows. The fisherman is doing applied wave physics with a stopwatch and a pole, and the bookkeeping closes without his ever being told the answer.*

HOT Example 11

A small loudspeaker radiates sound uniformly in all directions, with a total acoustic power output of 80W. Find (a) the intensity of the sound at a distance of 3.0m from the speaker, (b) the intensity at a distance of 6.0m from the speaker, (c) the ratio of the two intensities, and (d) the ratio of the corresponding amplitudes of the sound waves.

Solution

The speaker radiates uniformly, so the inverse-square law applies. At a distance r from a point source of total power P, the intensity is $I = P/(4\pi r^2)$.

(a) At $r_1 = 3.0m$:

$$I_1 = \frac{P}{4\pi r_1^{\,2}} = \frac{80W}{4\pi(3.0m^2)} = 0.707\ W/m^2$$

(b) At $r_2 = 6.0m$:

$$I_2 = \frac{P}{4\pi r_2^2} = \frac{80W}{4\pi(6.0m^2)} = 0.177\ W/m^2$$

(c) The ratio of intensities can be computed from the explicit values, or directly from the inverse-square law:

$$\frac{I_1}{I_2} = \left(\frac{r_2}{r_1}\right)^2 = \left(\frac{6.0m}{3.0m}\right)^2 = (2)^2 = 4$$

So the intensity at 3.0m is four times the intensity at 6.0m.

(d) The ratio of amplitudes follows from $I \propto A^2$. If the intensity ratio is 4, then the amplitude ratio is the square root of that:

$$\frac{A_1}{A_2} = \sqrt{\frac{I_1}{I_2}} = \sqrt{4} = 2$$

Equivalently, $A \propto 1/r$, so the amplitude at 3.0m is twice the amplitude at 6.0m.

Making sense of the answer: *The intensity at 3m is about 0.7 W/m², which is loud but realistic for a 1m-distance public-address speaker tested in an open field. By 6m the intensity has dropped fourfold to about 0.18 W/m², still loud but noticeably softer. A casual listener would describe the change as a reduction of about half in loudness, which on the decibel scale corresponds to roughly 6 dB, exactly the figure that an inverse-square doubling of distance produces. Everything is consistent.*

Think like a physicist: *The trap in this example is part (d). A student who has memorised "the intensity ratio is the inverse square of the distance ratio" may, while distracted, also write "the amplitude ratio is the inverse square of the distance ratio". This is wrong by a factor of 2 in the wrong direction. The correct rule is that amplitude scales as 1/r, not 1/r². The two rules differ by a square root because intensity is proportional to amplitude squared. Doubling the distance quarters the intensity but only halves the amplitude. Memorise the two scalings as a pair, and remember which one is which by the energy bookkeeping: it is energy that spreads over the spherical surface 4πr², so it is energy density (intensity) that falls as 1/r², and amplitude (which encodes the displacement, not the energy) follows the square root.*

Phase difference, path difference, and intensity. Three new ideas; one master equation $\Delta\varphi = (2\pi/\lambda)\Delta x$; one inverse-square law $I = P/(4\pi r^2)$; and one square-root subtlety that converts between intensity and amplitude. With these in hand, we are ready to ask the question that has been hovering in the background since the first stone hit Lake Victoria: *when two waves arrive at the same point, what does the medium actually do?* The next section answers that question with a single principle, the principle of superposition, and uses it to derive the constructive and destructive cases that are the bedrock of every interference calculation in physics. Phase difference is what we have just learned to compute. Superposition is what tells us what to do with the answer.

SUPERPOSITION OF WAVES

We have arrived at the most important principle in this chapter, and arguably in the next one as well. Phase difference, the subject of the previous section, gave us a way to compare two waves at two points. The principle of superposition, the subject of this section, tells us what to do with that comparison: *when two or more waves arrive at the same point at the same time, what does the medium actually do?* The answer is so simple that students sometimes mistrust it on first encounter, and so powerful that it underpins every interference, diffraction, and resonance calculation in the rest of the chapter and in the whole of optics in Chapter 8.

Mr. Akilikubwa illustrated the principle on Lake Victoria, by throwing a second stone into the water beside the first one. Where the two ripple patterns met, two of his observations are now ready to be promoted to working principles.

Observation one: the disturbances added at the meeting point. Where two crests crossed, a taller crest briefly stood up; where a crest met a trough, the water momentarily went flat.

Observation two: after the crossing, the two patterns continued outward unchanged. Each ring kept its shape and its speed; neither ring carried away any record of having met the other.

Both observations come from a single physical statement, which we are about to write down formally and use over and over again.

The principle of superposition

The principle states: *When two or more waves pass through the same point in a medium simultaneously, the resultant displacement of the medium at that point is the algebraic sum of the displacements that each wave would produce there individually.* Mathematically, if y_1 is the displacement that the first wave would produce on its own, y_2 the displacement the second wave would produce on its own, and so on, then the displacement actually produced by all the waves together is:

$$y = y_1 + y_2 + y_3 + \ldots$$

That is the entire principle. Three things about it are worth pinning down before we begin to use it.

First, *the addition is algebraic.* A positive displacement and a negative displacement of the same magnitude cancel; two positive displacements reinforce; two negative displacements also reinforce, in the negative direction. The principle does not say that crests add and troughs subtract; it says that displacements add. The signs sort themselves out automatically.

Second, *the principle holds only in linear media*; media whose response to a small disturbance is proportional to the size of the disturbance, the way a Hooke's-law spring responds linearly to a small displacement. Almost every wave you will study in this book lives in a linear medium: a stretched string at modest amplitude, a sound wave in air at speech-level intensities, a light wave in vacuum or glass. There are exceptions (for instance the shock wave at the front of a supersonic jet), and for them the simple addition does not work, but those exceptions are not in the syllabus and do not affect any calculation here.

Third, and remarkably, the principle implies that *waves do not interact with each other*. They only interact with the medium they are passing through. Two waves crossing the same point both deposit their displacements there, but neither wave is altered by the encounter. After the crossing, each wave continues exactly as it would have if the other had never existed. This is why the two ripple patterns on Lake Victoria walked away from one another unchanged: the water at the meeting point was bobbing up and down by the algebraic sum of the two ripples' separate demands, and once the two patterns had passed through, the two demands were no longer being made at the same point, and the water reverted to obeying each ripple independently.

Constructive and destructive interference

Apply the principle now to the simplest non-trivial case: two waves of the same frequency, the same amplitude, and the same direction of travel, arriving at the same point with a definite phase difference $\Delta\varphi$. Call them:

$$y_1 = A\sin(\omega t) \quad \text{and} \quad y_2 = A\sin(\omega t + \Delta\varphi)$$

Their sum, by superposition, is:

$$y = y_1 + y_2 = A\big(\sin(\omega t) + \sin(\omega t + \Delta\varphi)\big)$$

Apply the sum-to-product identity $\sin\alpha + \sin\beta = 2\sin((\alpha+\beta)/2)\cos((\alpha-\beta)/2)$ with $\alpha = \omega t$ and $\beta = \omega t + \Delta\varphi$. The half-sum is $\omega t + \Delta\varphi/2$, and the half-difference is $-\Delta\varphi/2$ (whose cosine is the same as $\cos(\Delta\varphi/2)$ because cosine is even). The result is

$$y = \left(2A\cos\left(\frac{\Delta\varphi}{2}\right)\right)\sin\left(\omega t + \frac{\Delta\varphi}{2}\right)$$

This is itself a sinusoidal wave at the same frequency ω. The new amplitude, sitting outside the sine, is the bracket

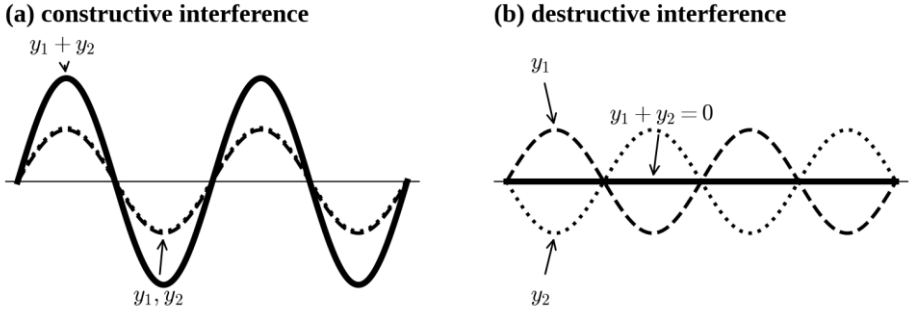
$$\text{resultant amplitude} = 2A\left|\cos\left(\frac{\Delta\varphi}{2}\right)\right|,$$

which depends only on $\Delta\varphi$. Two limiting cases dominate the rest of this chapter and all of Chapter 8.

Constructive interference: When $\Delta\varphi = 0, 2\pi, 4\pi, \ldots$ (an even multiple of π), the cosine is ± 1 and the resultant amplitude is $2A$, twice the amplitude of either wave alone. Two waves in phase, doing the same thing at the same time, reinforce each other completely. The medium at the meeting point oscillates with amplitude $2A$.

Destructive interference: When $\Delta\varphi = \pi, 3\pi, 5\pi, \ldots$ (an odd multiple of π), the cosine is 0 and the resultant amplitude is also 0. Two waves in antiphase, doing exact opposites, cancel each other completely. The medium at the meeting point does not oscillate at all.

For unequal amplitudes the same algebra runs through unchanged. The constructive case becomes $A_1 + A_2$ and the destructive case becomes $|A_1 - A_2|$. For equal amplitudes the constructive sum is $2A$ and the destructive sum is 0. The figure summarises both.



Conditions:

Constructive:	$\Delta\varphi = 0,\ 2\pi,\ 4\pi,\ \ldots$	resultant amplitude $= A_1 + A_2$
Destructive:	$\Delta\varphi = \pi,\ 3\pi,\ 5\pi,\ \ldots$	resultant amplitude $= \lvert A_1 - A_2 \rvert$

(For equal amplitudes, constructive gives $2A$ and destructive gives 0.)

Figure: ***Two-wave superposition.*** *(a) Constructive: two in-phase waves (dashed and dotted) of equal amplitude superpose to give a wave of double amplitude (solid). (b) Destructive: two antiphase waves (dashed and dotted) superpose to give zero (solid). The conditions on $\Delta\varphi$ for each case are stated below.*

What superposition does NOT do

It is worth a single paragraph to stress what the principle of superposition does **not** change.

- It does not destroy waves.
- It does not slow them down.
- It does not deflect them.

Two waves that meet at a point part company afterwards as if the meeting had never happened. This independence of waves is what makes interference patterns stable: *a single point in space, illuminated by two coherent light sources, settles into a definite condition (bright or dark or somewhere between) determined entirely by the phase difference at that point, and stays in that condition as long as the two sources keep their relative phase fixed.* Move the listener (or the screen) and the phase difference changes; move them back and the original condition returns. Superposition is reversible, geometrical, and humane: it adds; it never breaks.

Four worked examples follow.

BINDER Example 12

Two waves arriving at the same point in a medium are described by $y_1 = 5\sin(\omega t + 2\pi/3)$ m and $y_2 = 5\sin(\omega t - \pi/3)$ m. (a) Find the resultant displacement $y = y_1 + y_2$, and state the resultant amplitude. (b) Verify your part-(a) result independently by computing the phase difference $\Delta\varphi$ between the two waves and identifying the resulting interference type. (c) The amplitude of the second wave is now changed from 5m to 4m, with everything else (including its phase) unchanged. Find the new resultant amplitude.

Solution

(a) Apply the sum-to-product identity $\sin\alpha + \sin\beta = 2\sin((\alpha+\beta)/2)\cos((\alpha-\beta)/2)$ with $\alpha = \omega t + 2\pi/3$ and $\beta = \omega t - \pi/3$. The half-sum and half-difference are

$$\frac{\alpha + \beta}{2} = \omega t + \frac{\pi}{6}, \qquad \frac{\alpha - \beta}{2} = \frac{\pi}{2}$$

Putting this back into the identity:

$$y = 5\left(\sin\left(\omega t + \frac{2\pi}{3}\right) + \sin\left(\omega t - \frac{\pi}{3}\right)\right)$$

$$= 5 \times 2\sin\left(\omega t + \frac{\pi}{6}\right)\cos\left(\frac{\pi}{2}\right) = 10\sin\left(\omega t + \frac{\pi}{6}\right) \times 0 = 0$$

So the resultant displacement at every instant is zero. The resultant amplitude is 0.

(b) Compute the phase difference between the two given waves:

$$\Delta\varphi = \left(\omega t + \frac{2\pi}{3}\right) - \left(\omega t - \frac{\pi}{3}\right) = \frac{2\pi}{3} + \frac{\pi}{3} = \pi\,\text{rad}$$

A phase difference of π places the two waves in antiphase. The constructive-and-destructive bookkeeping of this section says that two equal-amplitude waves in antiphase produce a resultant amplitude of $|A - A| = 0$. This matches part (a) exactly. The two routes (direct algebraic summation and the phase-difference shortcut) agree, as they must.

(c) The phase difference is unchanged at $\Delta\varphi = \pi$ (we changed only the amplitude of the second wave, not its phase). The two waves are still in antiphase, so the resultant amplitude is $|A_1 - A_2|$:

$$\text{resultant amplitude} = (5\text{m} - 4\text{m}) = 1\text{m}$$

With the small amplitude mismatch, the cancellation is no longer perfect. The medium at the meeting point now oscillates with amplitude 1m, one-fifth of the larger of the two incoming waves.

Making sense of the answer: *Two non-trivial waves, each of amplitude 5m and oscillating at angular frequency ω, conspire to produce a perfectly motionless medium at the point where they meet. This is the most striking consequence of superposition. Energy is not destroyed (the two waves carry energy past the meeting point unchanged), but at this one point in space the medium is held perfectly still by two competing pulls of exactly equal magnitude and opposite direction. Part (c) then breaks the perfection: a 1m mismatch in amplitude leaves a 1m residual oscillation. Perfect cancellation requires both the antiphase condition **and** an exact amplitude match. The first condition is easy to satisfy in the laboratory; the second is hard.*

Think like a physicist: *The sum-to-product identity is doing real work in part (a), but you do not actually need it to get the answer. The phase-difference calculation in part (b) tells you $\Delta\varphi = \pi$ directly, and from that you know without further algebra that two equal-amplitude waves in antiphase cancel. The full sum-to-product computation is the rigorous way; the phase-difference recognition is the fast way. Practise both. In examination conditions the fast way will save minutes; in physics conditions, the rigorous way will protect you from arithmetic slips. Part (c) shows that the phase-difference shortcut also handles the unequal-amplitude case for free: the destructive-interference rule reads $|A_1 - A_2|$, no extra algebra required. The honest physicist owns both routes and uses each when it earns its keep.*

REAL Example 13

A pair of noise-cancelling headphones works by detecting incoming environmental noise with a small microphone, generating an electronic copy of the noise, inverting that copy in phase by π radians, and playing the inverted copy through the speaker into the listener's ear. The incoming noise and the cancelling wave then meet at the eardrum and add by superposition. Suppose the incoming noise on a city bus is well-modelled as a progressive wave of amplitude A = 0.040mm at the eardrum, and the cancelling wave generated by the headphones has amplitude A' = 0.035mm with a relative phase shift of exactly π radians.

(a) Using the principle of superposition, explain why a phase shift of π radians produces silence when the two amplitudes are equal.
(b) Find the residual amplitude at the eardrum when A = 0.040mm and A' = 0.035mm.
(c) Express the residual as a percentage of the incoming amplitude, and convert this to a percentage of the original intensity.

Solution

(a) Let the incoming noise at the eardrum be $y_1 = A\sin(\omega t)$. A wave with the same frequency but phase-shifted by π radians is $y_2 = A'\sin(\omega t + \pi)$. Using the trigonometric identity $\sin(\theta + \pi) = -\sin(\theta)$, the cancelling wave can be rewritten as $y_2 = -A'\sin(\omega t)$. By superposition, the total displacement of the eardrum is:

$$y_{\text{total}} = y_1 + y_2 = A\sin(\omega t) - A'\sin(\omega t) = (A - A')\sin(\omega t)$$

The resultant amplitude is therefore $|A - A'|$. When A = A' the difference is zero and the eardrum experiences no displacement at all, resulting in complete silence.

(b) With the given numbers,

$$\text{residual amplitude} = (0.040\text{mm} - 0.035\text{mm}) = 0.005\text{mm}$$

(c) Express the residual as a percentage of the incoming amplitude:

$$\text{residual fraction of A} = \frac{0.005\text{mm}}{0.040\text{mm}} = \frac{1}{8} = 12.5\%$$

Since intensity is proportional to amplitude squared, the residual intensity is the square of the residual amplitude fraction:

$$\text{residual fraction of I} = (0.125)^2 = 0.0156 = 1.56\%$$

So the headphones reduce the intensity of the incoming noise to about 1.6% of its original level (which corresponds to about 18 decibels of attenuation).

Making sense of the answer: *Real noise-cancelling headphones achieve roughly 20 to 30 decibels of attenuation in their target frequency band, so 18 decibels from this calculation is realistic. The lesson buried in the calculation is that perfect noise cancellation requires both a phase shift of exactly π* ***and*** *an amplitude match exactly equal to the incoming amplitude. The phase shift on its own is necessary but not sufficient. Engineers building these devices spend most of their effort on the amplitude-matching half of the problem, because the phase inversion is electronically easy but the amplitude match requires real-time measurement and adjustment as the user moves through changing acoustic environments.*

Think like a physicist: *This example is the simplest non-trivial application of superposition, and it captures the entire architecture of every interference calculation in this chapter and in Chapter 8. Two waves arrive at the same point. Each contributes its own sinusoidal displacement. The resultant is the algebraic sum. When the two waves are in antiphase and have equal amplitudes, the sum is zero. When they are in phase and have equal amplitudes, the sum is double. Everything else, including the loudness pattern in a concert hall and the bright fringes in Young's experiment, is just this same calculation done with different specific phase differences and amplitude ratios. Master this small example and you have the structural insight that drives the rest of the chapter and an entire chapter of optics in Chapter 8.*

REAL Example 14

Two identical loudspeakers, each emitting a pure sinusoidal tone of frequency 680Hz, are placed 3.0m apart on a school playing field at Miono Secondary, both connected to the same source so they emit in phase. Kipute walks slowly along a straight line parallel to the line joining the speakers, at a distance of 5.0m from each speaker. As she walks she hears the loudness of the tone alternate between distinct "loud" zones and distinct "quiet" zones.

(a) Using the principle of superposition, explain qualitatively why the loudness varies as she walks.
(b) At the midpoint between the two speakers (the point exactly equidistant from both), what does she hear, and why?
(c) She moves a small distance until the path difference from her position to the two speakers equals $\lambda/2$. What does she hear at that position, and why?

Solution

(a) Each speaker emits a sound wave of identical frequency and identical amplitude. The two waves arrive at the listener's ear at the same instant, but having travelled, in general, two different distances. The path difference Δx (the difference between those two distances) translates into a phase difference $\Delta\varphi = (2\pi/\lambda)\Delta x$ at the listener's ear. Where Δx is zero or a whole number of wavelengths, Δφ is an even multiple of π and the two waves arrive in phase: superposition gives constructive interference and the listener hears the tone loudly. Where Δx is half a wavelength or any odd multiple of half a wavelength, Δφ is an odd multiple of π and the two waves arrive in antiphase: superposition gives destructive interference and the listener hears the

tone faintly or not at all. As the student walks, her path difference changes continuously, and the loudness alternates between the constructive maxima and the destructive minima.

(b) At the midpoint, the listener is equidistant from both speakers, so the path difference is zero. The two waves arrive in phase, the resultant amplitude is the sum of their individual amplitudes (here 2A, since they are identical), and the loudness is at its constructive maximum. The midpoint is therefore one of the loudest positions on her path.

(c) When the path difference is $\lambda/2$, the phase difference is:

$$\Delta\varphi = \frac{2\pi}{\lambda} \times \frac{\lambda}{2} = \pi \text{ rad}$$

The two waves arrive in antiphase. With equal amplitudes, the resultant amplitude is zero and (in a perfectly absorbing field with no reflections) the listener hears silence. In a real outdoor field the cancellation is not quite perfect, because the two speakers are not point sources and the listener's ear is not a single point, but the loudness drops to a clear minimum that any human ear can pick out as a quiet zone.

Making sense of the answer: *This is exactly what happens in any open-air concert at the National Stadium in Dar es Salaam, in any school assembly with two PA speakers running the same microphone, and in every public-address system that has more than one speaker. The room (or field) divides itself into a pattern of loud and quiet zones, set by the geometry of the speakers and the wavelength of the sound. Concert hall designers spend serious engineering effort on placing speakers so that no listener sits permanently in a quiet zone for the dominant frequencies of speech and music.*

Think like a physicist: *Notice the architecture of the argument. The path-difference equation $\Delta\varphi = (2\pi/\lambda)\Delta x$, established in the previous section, is the engine that turns geometry (the listener's position) into phase difference. Then the principle of superposition, established in this section, turns phase difference into amplitude. The two equations chain together to predict, from the speaker spacing, the wavelength, and the listener's position, exactly how loud the tone will be at every point on the field. This same chain of reasoning underwrites the entire physics of interference, in any wave at any wavelength. If you can do the calculation here for sound, you can do it for light in Chapter 8.*

HOT Example 15

Two waves of identical amplitude A, identical angular frequency ω, and identical wave number k travel along the same string in opposite directions. The wave moving in the positive x-direction is $y_1 = A\sin(\omega t - kx)$; the wave moving in the negative x-direction is $y_2 = A\sin(\omega t + kx)$. (a) Find the resultant displacement $y = y_1 + y_2$. (b) Identify the type of wave that the resultant describes. (c) Find the displacement at the point $x = \lambda/6$ at the instant t when $\sin(\omega t) = 1$.

Solution

(a) Apply superposition. The resultant is the sum:

$$y = y_1 + y_2 = A\sin(\omega t - kx) + A\sin(\omega t + kx)$$

Use $\sin\alpha + \sin\beta = 2\sin((\alpha+\beta)/2)\cos((\alpha-\beta)/2)$ with $\alpha = \omega t - kx$ and $\beta = \omega t + kx$. The half-sum is ωt; the half-difference is −kx (whose cosine is cos(kx) because cosine is even). The resultant becomes:

$$y = 2A\sin(\omega t)\cos(kx) = (2A\cos kx)\sin\omega t$$

(b) The resultant is not a travelling wave. A travelling wave's argument always contains the combination $\omega t \pm kx$, with t and x bound together inside a single sine. Here the spatial dependence (cos kx) and the temporal dependence (sin ωt) appear in separate factors. The wave does not advance through the medium. Each particle of the medium oscillates in time at angular frequency ω, with an amplitude $2A\cos(kx)$ that depends only on its location x. This kind of wave is called a **standing wave** (or **stationary wave**). (It will be the subject of the next two sections of this chapter.)

(c) At $x = \lambda/6$, the spatial argument is:

$$kx = \frac{2\pi}{\lambda} \times \frac{\lambda}{6} = \frac{\pi}{3}$$

So $\cos(kx) = \cos(\pi/3) = 1/2$. At the instant when $\sin(\omega t) = 1$, the displacement at $x = \lambda/6$ is

$$y = \left(2A \times \frac{1}{2}\right) \times 1 = A$$

So the particle at $x = \lambda/6$ reaches a peak displacement of A, which is half the maximum standing-wave amplitude of 2A. It is not at a node (the cosine factor is non-zero there), and it is not at an antinode (the cosine factor is not ±1 there). It is at an intermediate amplitude.

Making sense of the answer: *The algebra has just produced a standing wave out of two travelling waves. Read this carefully: the principle of superposition, the sum-to-product identity, and the simple fact that two waves can travel in opposite directions on the same medium have together delivered the most physically interesting object in the chapter. Every guitar string, every organ pipe, every cantilever bridge swaying in the wind, every microwave oven cavity, and every electron's standing pattern around a hydrogen nucleus is a special case of this result. The next two sections take this expression apart and use it to derive the harmonics of stringed instruments and the boundary-condition rules that select which standing waves an instrument can support.*

Think like a physicist: *Notice the structural difference between (2A cos kx) sin ωt and Asin(ωt − kx). The second is a travelling wave: t and x appear together inside a single function. The first is a standing wave: t and x appear in separate functions, multiplied together. This separation of variables, in mathematics, is a deep structural fingerprint, and the physics fingerprint matches: the travelling-wave shape moves through the medium without changing, and the standing-wave shape stays in place while changing in amplitude. Two superficially similar equations describe two completely different physical objects. The arithmetic that produced the standing wave from two opposite-direction travelling waves is honest: the sum-to-product identity is the only step, and it is the algebraic shadow of the physical fact that two opposite-going waves on the same line can fold themselves into a single oscillation locked in place. Hold this calculation in your head; the next section's derivation of node and antinode positions falls out of it without further work.*

Superposition is now in our toolkit. We know how to add two waves at the same point, what the result looks like for special phase differences (constructive 2A, destructive 0), and what the result looks like when the two waves are travelling in opposite directions on the same medium (a standing wave). The next section uses superposition for a slightly different special case: two waves of nearly the same frequency, travelling in the same direction. The result there is not a steady interference pattern but a slow rhythmic rise and fall of loudness called beats, and the physics of beats is the foundation of every musician's tuning ear and every stationary-engine vibration calibration.

BEATS

In the previous section we used the principle of superposition on two waves of the same frequency. The two waves either reinforced (constructive) or cancelled (destructive), depending on the fixed phase difference between them, and the result was a steady pattern in space. We turn now to the slightly different and altogether more musical case: two waves of nearly but not quite the same frequency, travelling in the same direction, arriving at the same point. The result is no longer a steady pattern. It is *a periodic rise and fall of loudness* called **beats**, which every guitar tuner, piano technician, and physics student with two slightly mistuned notes has heard. Beats are how musicians tune by ear, how engineers detect tiny frequency differences in machinery, and how a Form Five student can learn, with no equipment more sophisticated than a tuning fork and a guitar string, that two frequencies separated by 1Hz announce themselves as one full cycle of loud-quiet-loud per second.

This section does three things. It describes what beats sound like, derives the beat-frequency formula in two complementary ways (algebraic sum-to-product, and direct cycle-counting), and sets up the worked examples that come at the end. The arithmetic is tiny. The physics is large.

What a beat actually is

Imagine two pure tones being sounded at the same point in your ear. Tone 1 has frequency f_1; tone 2 has frequency f_2; both have the same amplitude A; they arrive at the eardrum at the same time, and they superpose by the principle of the previous section. If f_1 and f_2 are very different (say one at 440Hz and another at 880Hz), you hear the two tones as separate notes; the brain has no trouble identifying them. But if f_1 and f_2 are very close (say one at 440Hz and one at 442Hz), the brain cannot resolve them. It interprets the superposition as a single tone whose loudness rises and falls with a slow rhythmic pulse. The **pulse rate**, which is *how many loudness peaks the ear hears per second*, is the **beat frequency**, written $\mathbf{f_b}$. For our 440 and 442 example, the beat frequency turns out to be 2Hz, meaning two loud-quiet-loud cycles per second. The result is the kind of slow throbbing that anyone who has tuned a guitar knows.

Beats are not an illusion. They are the algebraic consequence of the principle of superposition. The next subsection writes them down formally.

The sum-to-product derivation

Let the two arriving waves be

$$y_1 = A\sin(2\pi f_1 t) \quad \text{and} \quad y_2 = A\sin(2\pi f_2 t)$$

By superposition, the total displacement of the eardrum at time t is

$$y = y_1 + y_2 = A\big(\sin(2\pi f_1 t) + \sin(2\pi f_2 t)\big)$$

Apply the sum-to-product identity $\sin\alpha + \sin\beta = 2\sin\left(\frac{\alpha+\beta}{2}\right)\cos\left(\frac{\alpha-\beta}{2}\right)$ with $\alpha = 2\pi f_1 t$ and $\beta = 2\pi f_2 t$. The half-sum is $\pi(f_1 + f_2)t$ and the half-difference is $\pi(f_1 - f_2)t$. Substituting:

$$y = 2A\cos(\pi(f_1 - f_2)t)\sin(\pi(f_1 + f_2)t)$$

Read this expression with care, because it has two factors and they are doing two different jobs.

The second factor, $\sin(\pi(f_1 + f_2)t)$, is a fast oscillation at the average frequency $\frac{f_1+f_2}{2}$. For two close frequencies, this average is essentially the same as either f_1 or f_2. It is the tone the ear hears.

The first factor, $2A\cos(\pi(f_1 - f_2)t)$, is a slow oscillation at half the difference frequency. It plays the role of a **time-varying amplitude** or **envelope** for the fast oscillation. As time advances, the envelope swings between $+2A$ (when the cosine is $+1$, both waves are momentarily in phase, full constructive interference) and $-2A$ (when the cosine is -1, both waves are again momentarily in phase but on the negative half of the cycle), passing through zero (when the cosine is zero, the two waves are antiphase, full destructive interference) twice per envelope cycle.

Now the subtle step that produces the beat-frequency formula. The ear does not respond to the signed amplitude of the envelope. It responds to the loudness, which is the squared amplitude. In a single envelope cycle the cosine factor goes through the values $+2A$, 0, $-2A$, 0, $+2A$; the squared amplitude therefore goes through $4A^2$, 0, $4A^2$, 0, $4A^2$. Two loudness peaks per envelope cycle. Therefore the loudness pulse rate is twice the envelope frequency:

$$f_b = 2 \times \frac{|f_1 - f_2|}{2} = |f_1 - f_2|$$

This is the beat-frequency formula. **In words:** the **beat frequency** *is the absolute value of the difference between the two superposing frequencies.* Two notes one hertz apart give one beat per second. Ten hertz apart give ten beats per second. A hundred hertz apart give a hundred beats per second, which is too fast for the ear to resolve as beats and is heard instead as a third tone (the so-called difference tone of music theory).

The same result by counting cycles

There is a second derivation of the beat-frequency formula that uses no trigonometric identities, only the simple observation that two oscillators of different periods drift in and out of phase at a definite rate. Suppose, without loss of generality, that $f_1 > f_2$ and that the two waves happen to start in phase at $t = 0$. Wave 1 has period $T_1 = \frac{1}{f_1}$; wave 2 has period $T_2 = \frac{1}{f_2}$. Wave 1 oscillates a little faster than wave 2.

After some time T_b, the two waves are once again exactly in phase. During that interval, wave 1 has completed some integer number n of full cycles, and wave 2 has completed exactly $n - 1$ full cycles (it has fallen behind by one cycle). Both completed-cycle counts must equal the same time interval T_b:

$$T_b = nT_1 = (n - 1)T_2$$

Eliminate n by dividing the two equations. From the first, $n = \frac{T_b}{T_1}$; from the second, $n - 1 = \frac{T_b}{T_2}$. Subtracting,

$$\mathbf{1} = T_b\left(\frac{1}{T_1} - \frac{1}{T_2}\right) = \mathbf{T_b(f_1 - f_2)}$$

The reciprocal of T_b is the beat frequency:

$$f_b = \frac{1}{T_b} = f_1 - f_2$$

Restoring the absolute-value sign for the case $f_2 > f_1$ gives back the same formula as before. The two derivations agree, as they must: they are different paths through the same physics. The sum-to-product route gives you the envelope and the squared-amplitude reasoning that explains why the loudness pulses at the difference frequency rather than at half of it. The cycle-counting route gives you a quick mental shortcut: the two notes drift past each other at a rate of $|f_1 - f_2|$ cycles per second, and that drift rate **is** the beat rate.

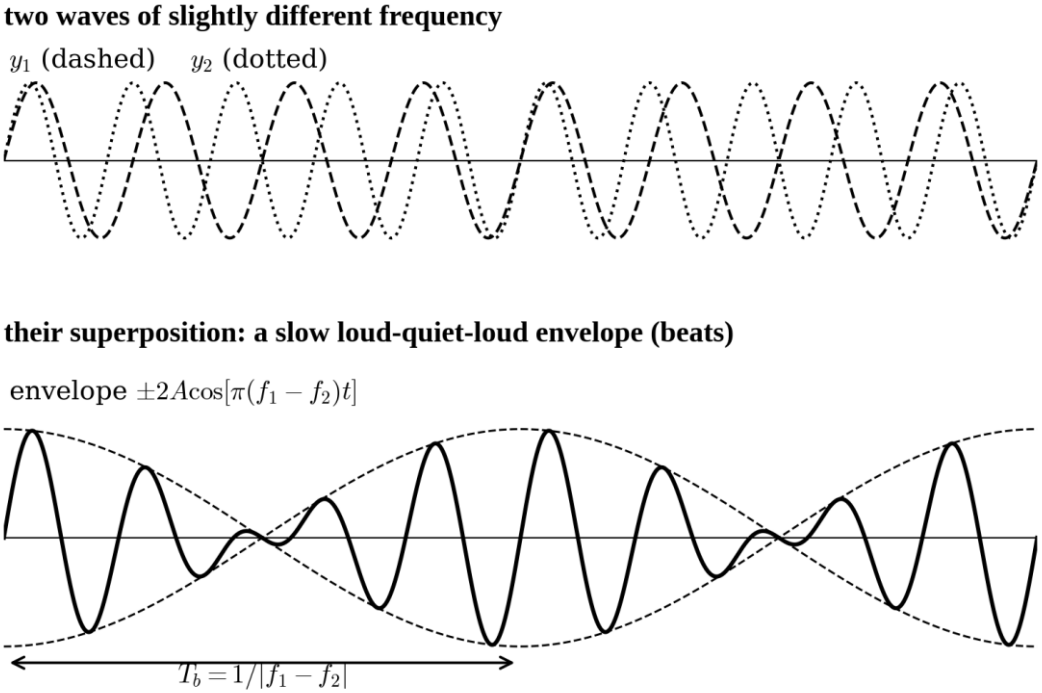


Figure: ***Beat formation.*** *Top: two waves of slightly different frequency, drawn dashed and dotted. Bottom: their superposition (solid), bounded by the slow envelope $\pm 2A\cos[\pi(f_1 - f_2)t]$ (dashed). One full beat period $T_b = 1/|f_1 - f_2|$ is marked. Within T_b the loudness rises from quiet to loud back through quiet.*

Before the ideas drift away, let us anchor them using two worked examples.

BINDER Example 16

Two tuning forks, P and Q, sounded together produce 6 beats per second. Fork P is in unison with a 40cm length of a sonometer wire, and fork Q is in unison with a 32cm length of the same wire under the same tension. Find the frequencies of the two forks.

Solution

The frequency of the fundamental on a sonometer wire under fixed tension is inversely proportional to the vibrating length: a wire half as long has twice the frequency. So with a fixed wire and fixed tension,

$$f \propto \frac{1}{L}$$

Fork P is in unison with the 40cm length and Fork Q with the 32cm length, so the ratio of their frequencies is the inverse ratio of their lengths:

$$\frac{f_P}{f_Q} = \frac{L_Q}{L_p} = \frac{32\text{cm}}{40\text{cm}} = \frac{4}{5}$$

So $f_P = \frac{4}{5} f_Q$. The shorter wire (32cm) sounds at the higher frequency, so $f_Q > f_P$. The beat-frequency formula combined with the unknown frequencies gives the second equation:

$$f_Q - f_P = 6\text{Hz}$$

Substitute $f_P = \frac{4}{5} f_Q$ into this:

$$f_Q - \frac{4}{5} f_Q = \frac{1}{5} f_Q = 6\text{Hz}$$

So $f_Q = 30\text{Hz}$, and $f_P = \frac{4}{5} \times 30\text{Hz} = 24\text{Hz}$.

Making sense of the answer: *Two unknown frequencies, two equations: one from the geometry of the sonometer (the ratio of the lengths) and one from the beat measurement (the difference of the frequencies). Two equations in two unknowns close the system without further information. The beat measurement is the part of the problem that closes it; without the 6Hz beats, the geometry alone would only tell us the ratio of the two frequencies, not the absolute values. This is exactly how a piano tuner converts the audible 6 beats per second between a tuning fork and a string into a numerical correction to the string's tension. The beats turn a comparison into a measurement.*

Think like a physicist: *Notice the structural lesson. Two pieces of physical information were given: the geometric relationship $f \propto 1/L$ (which lives in the chapter on stationary waves) and the beat-frequency relationship $f_b = |f_1 - f_2|$ (which is brand new in this section). Neither one on its own would have solved the problem; both together did. The chapter is now mature enough that examination questions will routinely require you to chain results from two or three of its sections in a single calculation. The discipline is to look at the given information and ask, before you reach for any equation: what physical relationship does this measurement give me, and which other section's relationship will I need to combine with it? Once you have that mental list of two or three equations, the algebra is mechanical.*

REAL Example 17

The school music room at Miono Secondary, Saturday afternoon. Kipanga is tuning his guitar's high E string against a 440Hz tuning fork, and Kipute and Mr. Akilikubwa have stopped by to listen. The fork is sounded; Kipanga plucks the string; the room fills with eight clear beats per second. Kipanga reaches for the tuning peg.

Kipanga: *"Eight beats. So the string is wrong by eight hertz. I know it is below the fork because the string sounds lower; let me tighten it."*

Kipute: *"Wait. You don't actually know it is below the fork. The beats only tell you the absolute difference."*

Kipanga: *"Then how do I know which way to turn the peg?"*

Kipute: *"You don't know yet. You will know in two seconds. Tighten the peg slightly. If the beats slow, you were below the fork and you are moving toward it. If the beats speed up, you were above and you are moving away from it. Either way, the string answers the question for you within one cycle of the beat."*

Mr. Akilikubwa: *"This is what physicists call breaking the absolute-value sign. The beat-frequency formula has a $|\cdot|$ in it, and the $|\cdot|$ hides one bit of information, the sign of $f_{string} - f_{fork}$. A small perturbation reveals that bit."*

Kipanga tightens the peg. The beats slow to six per second, then four, then two, then one, and then, at one particular angle of the peg, the room goes quiet. The string and the fork sing together as if they were the same instrument. Kipanga grins. Kipute nods. Mr. Akilikubwa goes back to his book.

From this dialogue, answer the following:

(a) From the original eight beats per second, find the two possible frequencies of the guitar string before any tightening took place.
(b) The beats slow as the string is tightened. From this single observation, determine which of the two possibilities in (a) was the actual starting frequency, and explain your reasoning physically.
(c) Find the frequency of the string at the moment of silence, and explain physically why "silence" is the equation $f_b = 0$ made audible.

Solution

(a) The beat-frequency formula gives the magnitude of the difference between the string's frequency and the fork's frequency:

$$f_b = |f_{string} - f_{fork}| = 8\text{Hz}$$

With $f_{fork} = 440\text{Hz}$, the absolute-value sign hides whether the string is below or above the fork. Both options satisfy the equation:

$$f_{string} = 440\text{Hz} \pm 8\text{Hz} = 432\text{Hz or } 448\text{Hz}$$

(b) Tightening a string raises its tension, which raises its frequency. So as Kipanga tightens, f_{string} increases. The observation is that the beats slow as he tightens, which means $|f_{string} - 440|$ is decreasing, which means f_{string} is approaching 440 from one side. There are two cases. If the string started at 432Hz, raising it moves it toward 440 and the beats decrease, consistent with what Kipanga heard. If the string started at 448Hz, raising it moves it away from 440 and the beats would increase, inconsistent with what he heard. Therefore the string was at 432Hz at the start.
(c) At the moment of silence, the beat frequency is zero:

$$f_b = |f_{string} - 440\text{Hz}| = 0$$

This requires $f_{string} = 440\text{Hz}$ exactly.
The string and the fork now sound at identical frequencies; they superpose constructively at every instant; they reinforce each other into a single steady tone with no envelope. The silence Kipanga heard is not absence of sound. It is absence of **variation** in the loudness. The note is still there, twice as loud as either source alone, but it does not pulse, because there is nothing for it to pulse against.

Making sense of the answer: *Every musician who has ever tuned by ear has solved exactly this physics problem, with no algebra and no awareness of the equation $f_b = |f_1 - f_2|$. They listen for the silence, and the silence is the equation. Kipute's argument in the dialogue is what physicists call a* ***perturbation method:*** *the ambiguity (above or below the fork?) cannot be resolved by listening, but it can be resolved by listening combined with one small action whose direction you control. The same logic underwrites every measurement of a tiny frequency shift, from the Doppler-shift radar speed gun to the gravitational redshift of an atomic clock at the top of a tall building. A small known perturbation breaks the absolute-value sign and tells you the sign of the unknown.*

Think like a physicist: *Mr. Akilikubwa's phrase "breaking the absolute-value sign" is worth keeping. Many physical measurements give you the magnitude of a signed quantity but not its sign. The beat-frequency formula is the textbook example. So is the Doppler shift formula for an unknown source moving along your line of sight. So is the phase difference at a single point in an interferometer. The measurement device hands you* $|x|$*; the missing bit of information is the* ***sign function*** $(\text{sgn}(x))$*; a deliberate perturbation reveals* $\text{sgn}(x)$ *within one beat period or one fringe shift. This is a pattern, and once you spot it you find it everywhere in experimental physics. The guitar tuner is doing it without knowing it is doing it. You, who have just read this dialogue, are doing it on purpose.*

Beats are now in our toolkit. Two waves of nearly equal frequency, superposing in the same direction, give a slow rhythmic envelope at the difference frequency, and that envelope is what the ear hears as beats. The next section moves to the third and final special case of superposition that this chapter cares about: two waves of **equal** frequency, equal amplitude, but travelling in **opposite** directions on the same medium. Their superposition will not produce a steady interference pattern, and it will not produce beats. It will produce something stranger and more important than either: a wave whose shape is fixed in space and whose amplitude oscillates in time. We have already seen the algebra (it appeared in HOT Example 15 of the previous section as a foreshadowing). The next section gives that algebra a name, a physical interpretation, and a string of consequences that govern every musical instrument in this chapter.

STATIONARY WAVES: FORMATION AND CHARACTERISTICS

In the previous section we superposed two waves of nearly equal frequency travelling in the same direction and got beats. We now do something seemingly minor and physically enormous: we superpose two waves of **exactly equal** frequency, equal amplitude, and equal wavelength, but travelling in **opposite** directions on the same medium. The two waves $y_1 = A\sin(\omega t - kx)$ and $y_2 = A\sin(\omega t + kx)$ superpose, and the sum-to-product identity gives:

$$y = \big(2A\sin(kx)\big)\cos(\omega t)$$

(or equivalently, with a different phase choice for the source, $y = \big(2A\cos(kx)\big)\sin(\omega t)$). This is not a wave that travels. It is a wave that stays in place while oscillating in amplitude. The shape is fixed in space; only the size of the displacement at each fixed point oscillates in time. This object is called a **stationary wave** or **standing wave**. This section gives that object

its physical interpretation, identifies the special points called **nodes** and **antinodes**, and lays out how a stationary wave differs from the progressive waves we have been studying so far.

The next section will use the same equation to compute tensions in strings. The section after that will derive the harmonic frequencies of every musical instrument made of a string fixed at both ends. Both of those calculations rest on the picture this section draws.

How a stationary wave actually forms in a string

The mathematics asks for two waves of identical frequency and amplitude travelling in opposite directions on the same medium. The physics offers two ways to arrange that.

The first way is to drive the medium from both ends with two synchronised oscillators of identical frequency and identical amplitude, the left source sending a wave rightward and the right source sending a wave leftward. The two waves meet in the middle and superpose. The arrangement is rare in practice (it requires very careful synchronisation of the two sources) and is mostly a theoretical idealisation.

The second way, which is far more common and is how every musical string in the universe actually does this, is to drive the medium from one end and let the wave reflect from the other end. A wave travelling rightward hits a fixed boundary, inverts in phase, and reflects back leftward. The reflected wave has the same frequency, the same amplitude (in an idealised medium with no dissipation), and the same wavelength as the incident wave; it just runs the other way. Incident and reflected superpose. If the source frequency is chosen so the boundary conditions are satisfied (we will sharpen this in the next section), the result is a stable stationary wave that rings indefinitely. Pluck a guitar string and what you have set up, in less than a millisecond, is exactly this two-wave reflection picture, with hundreds of round trips per second between the two fixed ends.

In either physical arrangement, the algebraic result is the same: a function of x multiplied by a function of t, separately. That separation of variables, with no $\omega t \pm kx$ appearing inside a single sine, is the algebraic fingerprint of a stationary wave. It is the physical fingerprint too: time and space are no longer bound together inside the wave's argument.

Nodes and antinodes: what they are and where they live

Read the equation $y = (2A\sin(kx))\cos(\omega t)$ carefully. The factor $2A\sin(kx)$ depends only on position. It tells us, at each x, how big the oscillation will be. The factor $\cos(\omega t)$ depends only on time. It tells us, at any given instant, what fraction of that maximum oscillation we are currently seeing. Two factors, two completely separate jobs.

Look first at the spatial factor. There are particular points on the string where $\sin(kx)$ is exactly zero. At those points, the displacement y is zero for all time, regardless of what the cosine factor is doing. These are the points where the string never moves. They are called **nodes** and we write them with the letter **N**. From $\sin(kx) = 0$, the nodes occur at $kx = 0, \pi, 2\pi, 3\pi, \ldots$, that is, at:

$$x = 0, \frac{\lambda}{2}, \lambda, \frac{3\lambda}{2}, 2\lambda, \ldots$$

Adjacent nodes are therefore $\frac{\lambda}{2}$ apart. Half a wavelength, not a whole one. **Memorise this.** *The space between two consecutive points where the string never moves is half the wavelength of the constituent travelling waves, not the full wavelength.*

Now look at the points where $\sin(kx)$ reaches its maximum magnitude of ± 1. At those points, the spatial factor multiplies the cosine by the largest possible amount; the displacement at those points oscillates between $+2A$ and $-2A$ as the cosine cycles. These are the points of maximum oscillation. They are called **antinodes** and we write them with the letter **A**. From $\sin(kx) = \pm 1$, the antinodes occur at $kx = \pi/2, 3\pi/2, 5\pi/2, \ldots$, that is, at:

$$x = \frac{\lambda}{4}, \frac{3\lambda}{4}, \frac{5\lambda}{4}, \ldots$$

Adjacent antinodes, like adjacent nodes, are $\frac{\lambda}{2}$ apart. The distance from a node to its nearest antinode is:

$$\frac{\lambda}{4}$$

a quarter of a wavelength. The N–A–N–A–N… pattern marches along the string with this fixed quarter-wavelength spacing, and the pattern is fixed: the same x position is a node now, ten seconds from now, and ten years from now, as long as the wave persists.

One important physical consequence of the node geometry. At a node, the medium never moves; therefore no energy is transported across a node, because energy transport requires both a force and a velocity, and a particle at a node has zero velocity at every instant. Energy in a stationary wave is stored locally, oscillating back and forth between kinetic (when the segment is moving past equilibrium) and potential (when the segment is at maximum displacement). It does not travel along the string the way energy in a progressive wave does. This is the most important physical difference between a stationary wave and a progressive wave, and it is the first item on the list that comes next.

How a stationary wave differs from a progressive wave

The algebra of a stationary wave is $y = (2A\sin(kx))\cos(\omega t)$. The algebra of a progressive wave is $y = A\sin(\omega t - kx)$. Five physical differences fall out of comparing those two equations, and each difference is a consequence of the algebra rather than something to be memorised in isolation. Read each as: see what the equation does, and read off the physics.

1. Energy

The progressive wave's argument $\omega t - kx$ advances steadily as the disturbance moves rightward, and energy is carried along with it; a wave train passing through a fixed point delivers a steady stream of energy per second. The stationary wave has no such directed transport. Its argument never combines t and x inside a single sine; the time and space dependences are factored apart. There is no propagation direction, and there is no net energy transport. Energy is stored, locally, between consecutive nodes.

2. Amplitude

In the progressive wave, every particle of the medium oscillates with the same amplitude A. The amplitude is constant in space. In the stationary wave, the amplitude depends on position: it is $|2A\sin(kx)|$, which varies from zero (at the nodes) to 2A (at the antinodes). Two particles in different parts of the same stationary wave oscillate at different sizes, and a third particle exactly at a node does not oscillate at all.

3. Phase

In the progressive wave, the phase of a particle depends continuously on its position; two particles at different x values are at different stages of the same cycle, and the phase difference between them grows linearly with their separation. In the stationary wave, all the particles between two adjacent nodes share the same phase: they all rise together to maximum displacement, all fall back through equilibrium together, all reach minimum together. The two segments on either side of a node are antiphase: while one segment is up, the other is down. There are exactly two phases in the entire stationary wave (separated by π), and every particle is in one of them.

4. Wavelength

In the progressive wave, the distance between consecutive crests, or between any two consecutive points doing the same thing, is the wavelength λ. In the stationary wave, the distance between consecutive nodes is $\frac{\lambda}{2}$, half the wavelength of the constituent travelling waves. A common student mistake is to read the visible spatial period of the standing pattern (one node-to-node distance) as the wavelength of the wave; it is not. The wavelength is twice that distance.

5. Waveform

In the progressive wave, the entire pattern of crests and troughs slides bodily along the medium at speed $v = \frac{\omega}{k}$. In the stationary wave, no part of the pattern moves. The shape of the string at three different times in the figure (a) $t = 0$, (b) $t = \frac{T}{4}$, and (c) $t = \frac{T}{2}$ looks the same in horizontal location but different in vertical scale: the nodes are always at the same x values, the antinodes are always at the same x values, and only the displacement at each fixed x rises and falls in time.

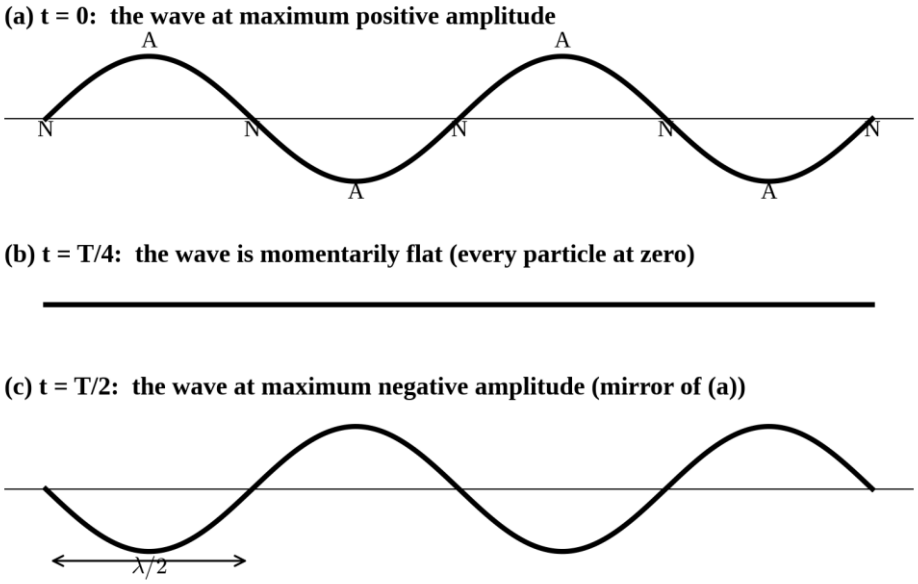


Figure: ***A stationary wave at three time instants.*** *(a) At t = 0 the wave is at maximum positive displacement; nodes (N) and antinodes (A) are labelled. (b) At t = T/4 every particle is momentarily at zero. (c) At t = T/2 the wave is at maximum negative displacement, the mirror of (a). The N and A positions are the same in all three panels: only the size of the displacement changes with time. The distance between two adjacent nodes is λ/2.*

With nodes and antinodes now simmering nicely on the string, let us serve them through two worked examples and enjoy the flavour of physics in action.

BINDER Example 18

A stationary wave is observed on a stretched string. Adjacent nodes on the string are 0.30m apart. A particle at the centre of one of the antinodes is observed to oscillate at frequency 50Hz, and the maximum displacement at that antinode is 4.0cm. Find: (a) the wavelength λ of the constituent travelling waves whose superposition produced this stationary wave, (b) the distance from any node to its nearest antinode, (c) the wave speed of the constituent travelling waves, and (d) the amplitude A of each of the two constituent travelling waves.

Solution

(a) The distance between adjacent nodes on a stationary wave equals $\frac{\lambda}{2}$. So:

$$\frac{\lambda}{2} = 0.30\text{m}, \quad \text{and} \quad \lambda = 2 \times 0.30\text{m} = 0.60\text{m}$$

(b) The distance from any node to its nearest antinode is half the node-to-node distance, equivalently a quarter of the wavelength:

$$\text{node} - \text{to} - \text{antinode distance} = \frac{\lambda}{4} = \frac{0.60\text{m}}{4} = 0.15\text{m}$$

(c) The wave speed comes from the standard relation $v = f\lambda$:

$$v = f\lambda = (50\text{Hz})(0.60\text{m}) = 30\text{m/s}$$

(d) The maximum displacement at any antinode of the stationary wave is 2A, where A is the amplitude of each of the two constituent travelling waves. So

$$2A = 4.0\text{cm}, \quad \text{and} \quad A = 2.0\text{cm}$$

Each of the two travelling waves whose superposition produced this stationary wave has an amplitude of 2.0cm. Together, when their crests align, they produce the 4.0cm maximum at the antinode.

Making sense of the answer: *All four numbers fall out of the stationary wave's geometry. The 0.30m node-to-node distance gives λ by doubling. The 0.15m node-to-antinode distance follows from λ by halving. The 30m/s wave speed follows from $v = f\lambda$. And the 2.0cm constituent amplitude follows from the antinode being twice as tall as a single travelling wave's crest. None of these calculations required the equation $y = (2A\sin(kx))\cos(\omega t)$ to be written down explicitly; the geometric facts of a stationary wave (adjacent nodes are $\frac{\lambda}{2}$ apart, antinodes oscillate at amplitude $2A$) were enough. The next section will use the equation directly to compute tensions; this section is the geometric warm-up.*

Think like a physicist: *Notice the structural lesson. A stationary wave is a wave whose properties have collapsed into spatial and temporal fingerprints that can be read separately. Geometry alone (without dynamics) gives you wavelength, node positions, antinode positions, and the constituent-wave amplitude. Dynamics enters through $v = f\lambda$ once a frequency is specified. The split is exactly the split between the two factors of the equation: $2A\sin(kx)$ lives in space and is set by the geometry of the boundary; $\cos(\omega t)$ lives in time and is set by whoever is driving the source. This factorisation is what makes a stationary wave such a useful structure for analysis: each factor can be examined and measured on its own, without the other.*

REAL Example 19

The physics laboratory at Miono Secondary, Wednesday afternoon. Mr. Akilikubwa has clamped a horizontal string of length 1.2m at both ends and connected one end to an electromagnetic vibrator. He has carefully tuned the driver frequency to one particular value at which the string settles into a single clean standing pattern: the two ends are perfectly still, the middle of the string sweeps up and down with a visible amplitude, and the rest of the string falls smoothly between those extremes. Kipanga and Kipute are watching.

Kipanga: *"The ends are not moving at all. The middle is moving a lot. Why does the string know where the ends are?"*

Kipute: *"It does not know anything. The ends are clamped. A clamped point has zero displacement always, and that is the same condition as a node. The string has no choice: at every instant the ends are at zero, so the standing pattern that lives on it must have nodes there. The middle is wherever the equation puts the next antinode after a node at one end."*

Kipanga: *"And the smooth shape between?"*

Kipute: *"Is $2A\sin(kx)$, with k set so that $\sin(kL) = 0$ too, since the far end is also a node. The boundary at one end is a forced condition the wave cannot escape. The boundary at the other end is the same condition, applied at $x = L$. Both ends being nodes selects which standing patterns can live on this string, and the simplest one is what you are looking at."*

Mr. Akilikubwa: *"Boundary conditions are the third partner in every standing-wave problem in this chapter. The wave equation provides the family of possible solutions; the boundary conditions select which member of the family actually shows up. Without boundary conditions a string would have no harmonics, no music, and no physics worth a Form Five examination paper."*

Use the demonstration to answer:

(a) The string is 1.2m long, fixed at both ends, and is observed to vibrate with one antinode at its midpoint and no other antinodes. Find the wavelength of the standing wave.
(b) Mr. Akilikubwa's driver is set to 50Hz. Find the speed of the transverse waves on the string.
(c) Mr. Akilikubwa now doubles the driver frequency to 100Hz. The string's tension and linear mass density are unchanged, so the wave speed is unchanged. Find the new wavelength of the standing wave on the string and the locations along the string of all nodes and antinodes in the new pattern.

Solution

(a) The two ends at $x = 0$ and $x = L = 1.2\text{m}$ are nodes. The midpoint at $x = \frac{L}{2} = 0.6\text{m}$ is an antinode. The distance from a node to the nearest antinode is $\frac{\lambda}{4}$, so:

$$\frac{\lambda}{4} = 0.6\text{m}, \quad \text{and} \quad \lambda = 2.4\text{m}$$

(b) From $v = f\lambda$:

$$v = (50\text{Hz})(2.4\text{m}) = 120\text{m/s}$$

(c) At 100Hz with $v = 120\text{m/s}$ unchanged, the new wavelength is

$$\lambda' = \frac{v}{f'} = \frac{120\text{m/s}}{100\text{Hz}} = 1.2\text{m}$$

The boundary condition still requires nodes at both ends. Adjacent nodes are $\frac{\lambda'}{2} = 0.6\text{m}$ apart. So in a 1.2m string we fit exactly two such intervals, which means three nodes in total: at $x = 0$, $x = 0.6\text{m}$, and $x = 1.2\text{m}$. The antinodes lie midway between consecutive nodes, so they are at $x = 0.3\text{m}$ and $x = 0.9\text{m}$. The pattern is N A N A N, with two antinodes instead of the previous one.

Making sense of the answer: *Doubling the driver frequency, with the wave speed unchanged, halved the wavelength and doubled the number of antinodes that fit between the two fixed ends. The string now hosts a more complex standing pattern, but the rule that selects which patterns can live there is unchanged: the two ends are forced to be nodes, and any wavelength compatible with that constraint is permitted. The smallest permitted wavelength fits one antinode (part a). The next smallest fits two antinodes (part c). The pattern continues: the next mode would fit three antinodes at* $\lambda'' = \frac{2L}{3} = 0.8\text{m}$*, at frequency 150Hz; the mode after that fits four, at 200Hz; and so on. We will systematise this counting in the next sections of the chapter.*

Think like a physicist: *The deep lesson is Mr. Akilikubwa's. The wave equation has a family of solutions, all of the form* $y = (2A\sin(kx))\cos(\omega t)$*, parameterised by which* k *you pick. Without boundary conditions, every* k *would be allowed and the string could ring at any frequency. The boundary conditions select the discrete set of allowed* k *values for which the chosen ends are nodes. That selection is what turns a continuum of possibilities into a discrete spectrum of harmonics, and that discrete spectrum is what gives a guitar a recognisable tone, a zeze a recognisable tone, and your own larynx a definite vowel rather than a continuous noise. Boundary conditions, applied to the wave equation, are the reason music exists.*

The picture is now in place. Stationary waves form when two opposite-direction travelling waves of equal frequency and amplitude superpose on the same medium; their displacement is $y = (2A\sin(kx))\cos(\omega t)$, with nodes spaced $\frac{\lambda}{2}$ apart and antinodes a quarter of a wavelength from each node. Energy is stored, not transported. Five differences from progressive waves emerge as direct consequences of the equation. The next section makes that equation the protagonist: it derives both the cosine form and the sine form, explains which form a particular boundary condition selects, and shows how to extract the wavelength, frequency, wave speed, and tension of a stretched string directly from a given stationary wave equation. The arithmetic ahead is short. The bookkeeping is exact.

THE STATIONARY WAVE EQUATION

The previous section drew the picture. Two travelling waves of equal frequency, equal amplitude, and equal wavelength pass through one another in opposite directions on the same medium, and the result is a wave whose shape stays put while every point oscillates in time. Nodes are the points that never move. Antinodes are the points that move by twice the constituent amplitude. The picture is complete; the equation is now ready to do work.

This section makes the equation the protagonist. It derives the equation algebraically from the superposition of two oppositely directed travelling waves, presents the two equivalent forms a stationary wave can take, shows how to read all the wave parameters off a given equation, and culminates in the calculation of greatest practical importance: extracting the tension of a stretched string directly from its stationary wave equation. The arithmetic ahead is straightforward. What matters is the order of the steps and what each step extracts.

Deriving the equation by superposition

Consider two transverse travelling waves of identical amplitude A, identical angular frequency ω, and identical wavenumber k, but moving in opposite directions along the same string. Take the first wave to be moving in the $+x$ direction and the second in the $-x$ direction:

$$y_1 = A\sin(\omega t - kx), \quad y_2 = A\sin(\omega t + kx)$$

By the principle of superposition the resultant displacement is $y = y_1 + y_2$, that is:

$$y = A\sin(\omega t - kx) + A\sin(\omega t + kx)$$

The two sines on the right have the same amplitude and add to a product by the standard sum-to-product identity:

$$\sin\alpha + \sin\beta = 2\sin\left(\frac{\alpha + \beta}{2}\right)\cos\left(\frac{\alpha - \beta}{2}\right)$$

With $\alpha = \omega t - kx$ and $\beta = \omega t + kx$, the half-sum is $\frac{\alpha+\beta}{2} = \omega t$ and the half-difference is $\frac{\alpha-\beta}{2} = -kx$. Substitute these gives:

$$y = 2A\sin(\omega t)\cos(-kx)$$

The cosine is an even function, so $\cos(-kx) = \cos(kx)$. The result is the equation we will live with for the rest of the chapter:

$$\mathbf{y = (2A\cos(kx))\sin(\omega t)}$$

Two factors. The factor $2A\cos(kx)$ depends only on position and not on time. The factor $\sin(\omega t)$ depends only on time and not on position. The product separates space from time completely, and that separation is the algebraic fingerprint of a stationary wave. No $\omega t - kx$ combination appears anywhere; nothing travels.

Two equivalent forms: which one shows up depends on the boundary

The derivation above started with sources whose phases were chosen so that at $t = 0$ both waves were in their pure Asin form. That phase choice produced the cosine spatial factor $2A\cos(kx)$ and the sine temporal factor $\sin(\omega t)$. A different phase choice for the same physical setup produces the equally valid form:

$$\mathbf{y = (2A\sin(kx))\cos(\omega t)}$$

Both forms describe stationary waves. They differ only in where the nodes sit. In the cosine spatial form, $\cos(kx)$ is zero where $kx = \frac{\pi}{2}, \frac{3\pi}{2}, \frac{5\pi}{2}, \ldots$, that is at $x = \frac{\lambda}{4}, \frac{3\lambda}{4}, \frac{5\lambda}{4}, \ldots$. The **point $x = 0$ is an antinode** in this form. In the sine spatial form, $\sin(kx)$ is zero at $x = 0, \frac{\lambda}{2}, \lambda, \frac{3\lambda}{2}, \ldots$, and the **point $x = 0$ is itself a node.**

Which form actually shows up on a real string is decided by the physical boundary at $x = 0$. A string clamped at $x = 0$ cannot move there, so its displacement is zero at that point at every instant. The only form that satisfies this is the sine spatial form, $y = (2A\sin(kx))\cos(\omega t)$, which gives $y(0, t) = 0$ identically. A string with a free end at $x = 0$ (one end loose, the other clamped, as in some pipe-end analogues we will meet in Chapter 9) would by contrast select the cosine spatial form, which puts an antinode at the free end. The mathematics offers both forms; the boundary condition picks the one that fits.

Reading the parameters off a given equation

Whenever a stationary wave equation is given to us, the procedure for extracting the wave parameters is mechanical. Suppose we are handed:

$$y(x, t) = (2A\sin(kx))\cos(\omega t)$$

with specific numbers in place of 2A, k, and ω. The four parameters fall out in a fixed order:

1) **Read the wavenumber k off the spatial argument and convert to wavelength.** The number multiplying x inside the trigonometric function is k, in radians per metre. The wavelength is:
$$\lambda = \frac{2\pi}{k}$$
2) **Read the angular frequency ω off the temporal argument and convert to frequency.** The number multiplying t inside the trig function is ω, in radians per second. The frequency is:
$$f = \frac{\omega}{2\pi}$$
3) **Compute the wave speed of the constituent travelling waves.** Either $v = f\lambda$ or, more directly,
$$v = \frac{\omega}{k}$$
4) **Read the constituent amplitude A off the leading factor.** The maximum amplitude at any antinode is the leading factor 2A, so the amplitude of each of the two travelling waves whose superposition produced this stationary wave is half of that, which is A.

Finding the tension of a stretched string from the equation

In Chapter 6 (simple harmonic motion of a stretched string), we established that the speed of a transverse wave on a stretched string of linear mass density μ under tension T is:

$$v = \sqrt{\frac{T}{\mu}}$$

Square both sides:

$$v^2 = \frac{T}{\mu}, \quad \text{so} \quad T = \mu v^2$$

This is the closing step of the whole procedure. Once we have read k and ω off the equation and computed $v = \omega/k$, the tension follows from $T = \mu v^2$ provided the linear mass density μ is known. If only the total mass m and the length L of the string are given, the linear mass density is $\mu = \frac{m}{L}$, and the calculation continues as before.

That is the whole arithmetic of this section, end to end:

$$\text{equation} \rightarrow k \rightarrow \lambda \rightarrow \omega \rightarrow f \rightarrow v = \frac{\omega}{k} \rightarrow T = \mu v^2$$

Read the wavenumber, read the angular frequency, divide them to get the wave speed, multiply the speed squared by the linear mass density. Six lines, one tension. The worked examples below make the procedure automatic.

So before the spatial and temporal factors of the equation tangle any further, let us straighten them out using a few worked examples.

BINDER Example 20

A stationary wave on a long stretched string has the equation

$$y(x,t) = 6\cos(3x)\sin(4t) \quad \text{cm},$$

with x measured in metres and t in seconds. (a) State the type of wave the equation describes and give the algebraic feature of the equation that confirms your identification. (b) Find the wavelength, frequency, wave speed of the constituent travelling waves, and the amplitude of each of those constituent travelling waves. (c) Find the positions of the first four nodes of the stationary wave on the positive x axis. (d) Write down the equations of the two component progressive waves whose superposition produced this stationary wave.

Solution

(a) **Stationary wave.** The equation has the form $y = \big(2A\cos(kx)\big)\sin(\omega t)$. The position factor $\cos(3x)$ and the time factor $\sin(4t)$ are multiplied together rather than appearing as a single combined argument $\omega t \pm kx$. This product structure is the algebraic fingerprint of a **stationary wave**. Equivalently, points x where $\cos(3x) = 0$ remain at zero displacement for all t: these are the nodes, and their fixed positions in space confirm that nothing is travelling.

(b) Read the parameters off the equation in the standard order.
Wavenumber: $k = 3\ \text{rad/m}$, so:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{3}\text{m} \approx 2.09\text{m}$$

Angular frequency: $\omega = 4\ \text{rad/s}$, so:

$$f = \frac{\omega}{2\pi} = \frac{4}{2\pi}\text{Hz} = \frac{2}{\pi}\text{Hz} \approx 0.637\text{Hz}.$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{4}{3}\text{m/s} \approx 1.33\text{m/s}.$$

Constituent amplitude: the leading factor is $2A = 6.0\text{cm}$, so:

$$A = 3.0\text{cm}.$$

(c) The nodes are the points where the position factor vanishes:

$$\cos(3x) = 0 \quad \Rightarrow \quad 3x = \frac{\pi}{2}, \frac{3\pi}{2}, \frac{5\pi}{2}, \frac{7\pi}{2}, \ldots$$

Solve for x and keep the first four positive values:

$$x = \frac{\pi}{6}\text{m}, \frac{\pi}{2}\text{m}, \frac{5\pi}{6}\text{m}, \frac{7\pi}{6}\text{m} \quad \approx 0.524\text{m}, 1.571\text{m}, 2.618\text{m}, 3.665\text{m}.$$

Adjacent nodes are $\frac{\pi}{3}\text{m} \approx 1.047\text{m}$ apart. This is exactly $\frac{\lambda}{2}$, as it must be for a stationary wave.

(d) The two component progressive waves are obtained by undoing the sum-to-product identity. Both travel with amplitude $A = 3.0\text{cm}$, angular frequency $\omega = 4\ \text{rad/s}$, and wavenumber $k = 3\ \text{rad/m}$, in opposite directions:

$$y_1 = 3\sin(4t - 3x)\text{cm}, \quad y_2 = 3\sin(4t + 3x)\text{cm}.$$

Their sum is exactly the given stationary wave equation, as the derivation at the start of the section confirms.

Making sense of the answer: *Every parameter of the stationary wave was extracted by inspection of two coefficients,* k *and* ω*. The amplitude required reading the leading factor and dividing by two. The node positions required nothing but solving* $\cos(kx) = 0$*. The component travelling waves required only writing back what the sum-to-product identity had hidden. The*

lesson is that a stationary wave equation hides nothing once you know which coefficient lives in space and which lives in time. The factor $\cos(3x)$ *is the spatial fingerprint, locking the node positions in place at* $x = \frac{\pi}{6}\text{m}, \frac{\pi}{2}\text{m}, \ldots$ *The factor* $\sin(4t)$ *is the temporal fingerprint, telling every non-node point of the string to oscillate at* $\omega = 4\,\text{rad/s}$*. The two fingerprints share no variable, and that is what lets us read them independently.*

Think like a physicist: *Note carefully what the cosine spatial form is telling us about the boundary at* $x = 0$*. At* $x = 0$ *the equation gives* $y(0,t) = 6.0\sin(4t)\text{cm}$*: the point at the origin oscillates with the maximum amplitude. The origin of this string is an antinode, not a node. If the string were clamped at the origin we would not be looking at this equation; we would be looking at the sine spatial form. The form of the equation, in other words, is itself a piece of physical information about how the string is supported at* $x = 0$*. Whenever you meet a stationary wave equation, the very first reflex should be to ask which form it is in and what that form is telling you about the boundary. The arithmetic comes second; the boundary lives in the equation itself.*

REAL Example 21

On the Kigamboni Bridge in Dar es Salaam, an engineer has fitted a strain-gauge sensor to one of the steel stay cables. After the morning rush of cargo lorries crosses the deck, the sensor records a transverse vibration on a particular cable described by:

$$y(x,t) = 0.8\cos\left(\frac{\pi x}{6}\right)\sin(40\pi t) \quad \text{cm},$$

with x measured along the cable in metres from the deck-end anchor and t in seconds. The cable's linear mass density is $\mu = 12\,\text{kg/m}$ and its total free length between the two anchors is $L = 18\text{m}$.

(a) Identify the type of wave the equation describes and state two algebraic features of the equation that confirm your identification.
(b) Find the wavelength, frequency, and wave speed of the transverse vibration on the cable.
(c) Find the tension in the cable.
(d) The cable is 18m long. How many half-wavelengths fit between the two anchors? Examine whether the equation's prediction at $x = 0$ is consistent with that end being a rigid anchor, and explain what the engineer should conclude.

Solution

(a) The equation has the form $y = (2A\cos(kx))\sin(\omega t)$. Two features confirm that this is a **stationary wave**.

First, the position factor and the time factor appear as a product, not as a single combined argument $\omega t \pm kx$: there is no travelling phase.

Second, the equation predicts fixed points (nodes) where $\cos\left(\frac{\pi x}{6}\right) = 0$, which remain at zero displacement for all t: such fixed nodes cannot exist on a progressive wave.

(b) Read the parameters off in the usual order.
Wavenumber: $k = \frac{\pi}{6}\text{rad/m}$, so:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{\pi/6}\text{m} = 12\text{m}$$

Angular frequency: $\omega = 40\pi\,\text{rad/s}$, so:

$$f = \frac{\omega}{2\pi} = \frac{40\pi}{2\pi}\text{Hz} = 20\text{Hz}.$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{40\pi}{\pi/6}\text{m/s} = 240\text{m/s}$$

(c) The cable obeys $v = \sqrt{\frac{T}{\mu}}$, so $T = \mu v^2$:

$$T = (12\,\text{kg/m})(240\,\text{m/s}^2)^2 = 6.91 \times 10^5\text{N}$$

(d) Compare the cable length to the wavelength:

$$\frac{L}{\lambda/2} = \frac{18\text{m}}{6\text{m}} = 3.$$

Three half-wavelengths fit exactly between the two anchors.

Now examine the boundary at $x = 0$. The equation gives $y(0,t) = 0.8\cos(0)\sin(40\pi t) = 0.8\sin(40\pi t)\text{cm}$: the point at the deck-end anchor oscillates with the full antinode amplitude of 0.8cm. But a properly anchored stay cable should have zero displacement at its anchor: that point is fixed to the deck and cannot move. The equation predicts an antinode at the anchor, while the physical setup demands a node there. The two are inconsistent.

The engineer should conclude that the deck-end anchor is not behaving as a rigid fix during the recorded vibration. There is play at that anchor, perhaps a worn bushing or a loosened socket allowing the cable end to move freely. The cable's vibration mode has therefore become close to that of a free-fixed system rather than the design-intended fixed-fixed system, which is a structural concern requiring inspection of the deck-end anchor before the next rush of traffic.

Making sense of the answer: *The arithmetic on this cable runs in five lines: $k \to \lambda, \omega \to f, v = \omega/k, T = \mu v^2$. Each line takes one substitution. The numerical answer $T = 6.91 \times 10^5 N$ is large, but it is the correct order of magnitude for a major stay cable carrying a span of bridge deck. Where the worked example earns its keep is part (d). The equation form has built-in information about the boundaries, and reading that information off the equation is part of what makes the calculation engineering rather than algebra. An antinode at $x = 0$ instead of a node is not a mathematical curiosity; it is a measurement that asks for action, because rigid anchors and antinodes are not compatible. A bridge engineer who does not check this part is just doing arithmetic.*

Think like a physicist: *There is a deeper lesson hiding in this example. The equation form (sine versus cosine in the spatial factor) is itself a measurement. The numerical coefficients k, ω, and the leading amplitude tell us the wavelength, frequency, speed, and tension. But which trigonometric function appears in the spatial factor tells us the boundary condition that the wave is satisfying at $x = 0$. A stationary wave on a fixed-end cable should be in the sine spatial form; the cosine spatial form means the cable is behaving as if it had a free end at $x = 0$. When the form contradicts the design, the design is wrong, or the cable is wrong, or the anchor is wrong. The equation cannot lie, because the equation came from the sensor; the only thing left to question is the cable's claimed boundary condition. This is how a physicist or an engineer reads a stationary wave equation: the parameters are arithmetic, but the form is diagnostic.*

HOT Example 22

A stretched string of length $\mathrm{L} = 1.5\mathrm{m}$ and mass $\mathrm{m} = 3.0 \times 10^{-2}$ kg carries a stationary wave described by:

$$\mathrm{y(x,t)} = 0.06 \sin\left(\frac{2\pi \mathrm{x}}{3}\right) \cos(120\pi \mathrm{t}) \quad \mathrm{m},$$

with x in metres and t in seconds. (a) Find the wavelength, frequency, and wave speed of the transverse vibration on the string. (b) Calculate the tension in the string. (c) Find which harmonic of the string is excited.

Solution

(a) Wavenumber: $\mathrm{k} = \frac{2\pi}{3}\mathrm{rad/m}$, so

$$\lambda = \frac{2\pi}{\mathrm{k}} = \frac{2\pi}{2\pi/3}\mathrm{m} = 3\mathrm{m}.$$

Angular frequency: $\omega = 120\pi\ \mathrm{rad/s}$, so

$$\mathrm{f} = \frac{\omega}{2\pi} = \frac{120\pi}{2\pi}\mathrm{Hz} = 60\mathrm{Hz}.$$

Wave speed:

$$\mathrm{v} = \frac{\omega}{\mathrm{k}} = \frac{120\pi}{2\pi/3}\mathrm{m/s} = 180\mathrm{m/s}.$$

(b) The linear mass density of the string is

$$\mu = \frac{\mathrm{m}}{\mathrm{L}} = \frac{3.0 \times 10^{-2}\ \mathrm{kg}}{1.5\mathrm{m}} = 0.020\ \mathrm{kg/m}.$$

From $\mathrm{v} = \sqrt{\frac{\mathrm{T}}{\mu}}$ we have $\mathrm{T} = \mu \mathrm{v}^2$:

$$\mathrm{T} = (0.020\ \mathrm{kg/m})(180\ \mathrm{m/s^2})^{\,2} = 648\mathrm{N}.$$

(c) The string is fixed at both ends, so both ends must be nodes. The sine spatial form $\sin\left(\frac{2\pi \mathrm{x}}{3}\right)$ puts a node at $\mathrm{x} = 0$. At $\mathrm{x} = \mathrm{L} = 1.5\mathrm{m}$:

$$\sin\left(\frac{2\pi \times 1.5}{3}\right) = \sin(\pi) = 0,$$

so the far end is also a node, exactly as required. The number of half-wavelengths between the two ends is

$$\mathrm{n} = \frac{\mathrm{L}}{\lambda/2} = \frac{1.5\mathrm{m}}{1.5\mathrm{m}} = 1,$$

so the string is vibrating in its **first harmonic**, that is, the fundamental mode.

Making sense of the answer: *The chain $k \to \lambda \to \omega \to f \to v = \omega/k \to T = \mu v^2$ runs the question from start to finish in six lines, and every line is a single substitution. The wave speed $v = 180 m/s$ is well within the range expected for a stretched musical string under tension. The tension $T = 648 N$ corresponds to about 66 kg of pull on a 1.5m string of mass 30g, large enough that the string is plausibly a steel piano-style wire rather than a guitar nylon. The fundamental wavelength on a 1.5m fixed-fixed string is $\lambda = 2L = 3m$, which matches the value the equation gave us; the equation and the boundary*

conditions agree on the mode. That self-consistency check is what tells us we have read the equation correctly: a fixed-fixed string of length $L = 1.5m$ cannot host a wavelength of 3m in any mode other than the first harmonic.

Think like a physicist: *This is exactly how a stationary wave equation wants to be read. It hands you k and ω; you produce λ, f, and v by elementary division; the string's mass and length give you μ; the formula $T = \mu v^2$ delivers the tension. The harmonic is identified by counting half-wavelengths against the string length. Six steps. No memorisation, no tricks, no hidden constants. The reason this question is HOT is not that any single step is hard but that all six must be performed in order, with the right physical relation invoked at each step. Mastering this sequence by repetition is what we have done in this section, and what every well-set examination paper will continue to test.*

The equation $y = (2A\sin(kx))\cos(\omega t)$ is now machinery. We have derived it by superposition; we have shown that the cosine and sine spatial forms are equivalent up to a phase, with the boundary condition selecting between them; we have read off the wavelength, frequency, wave speed, and constituent amplitude from a given equation in a fixed order; and we have closed the loop on stretched strings by computing the tension from the equation and the linear mass density. The next section takes the same equation, applies the boundary condition that the two ends of a string are fixed, and discovers the harmonic spectrum that gives every stringed musical instrument its identity. The arithmetic that follows will rest squarely on the procedure made automatic in this section.

VIBRATIONS OF STRINGS: HARMONICS AND THE SONOMETER

The previous section showed how to read every parameter of a stationary wave off its equation, and how to compute the tension of a stretched string from those parameters. We turn now to the question that turned the stretched string into the foundation of musical instruments and acoustic measurement: *what frequencies, exactly, can a string fixed at both ends be made to vibrate at?* The answer is one of the most elegant in classical physics. Out of a continuum of mathematical possibilities, the boundary condition that the two ends are fixed selects only a discrete sequence of frequencies. Each frequency is an integer multiple of a single fundamental, and the spacing between successive allowed frequencies is the fundamental itself.

Three short calculations make this discrete spectrum explicit. From the spectrum follow three universal laws of stringed instruments, expressing how the fundamental depends on length, tension, and linear mass density. Every guitar, *kora*, *zeze*, violin, and piano on Earth obeys those three laws to the letter; the same laws govern the sonometer in the laboratory. Once the laws are written down, the rest of the section is exercise: given any three of the four quantities (length, tension, linear mass density, harmonic frequency), the fourth follows from a single line of arithmetic.

From the boundary condition to the harmonic spectrum

Take a uniform stretched string of length L, clamped rigidly at both ends. The two fixed ends are at $x = 0$ and $x = L$. From the previous section, the stationary wave equation that has a node at $x = 0$ is the sine spatial form:

$$y(x,t) = (2A\sin(kx))\cos(\omega t)$$

The boundary condition at the far end requires that $y(L,t) = 0$ for all t, which forces the spatial factor to vanish at $x = L$:

$$\sin(kL) = 0.$$

The sine function is zero only at integer multiples of π, so the boundary condition selects

$$kL = n\pi, \quad n = 1, 2, 3, \ldots$$

The integer n is called the **harmonic number**. The case $n = 0$ is excluded because it gives $k = 0$ and a zero wave. Negative integers give the same physical solutions as their positive counterparts, and so are not listed separately. The wavenumber that is allowed in the n-th harmonic is therefore

$$k_n = \frac{n\pi}{L}, \quad \text{so} \quad \boldsymbol{\lambda_n = \frac{2\pi}{k_n} = \frac{2L}{n}} \text{ or } \mathbf{n = \frac{L}{\left(\frac{\lambda_n}{2}\right)}}$$

This is the central geometric statement of the section: *in the n-th harmonic, exactly n half-wavelengths fit between the two clamped ends.* The fundamental ($n = 1$) fits one half-wavelength, so $\lambda_1 = 2L$. The second harmonic ($n = 2$) fits two half-wavelengths, so $\lambda_2 = L$. The third harmonic ($n = 3$) fits three, so $\lambda_3 = 2L/3$. And so on. The following figure, shows the first three.

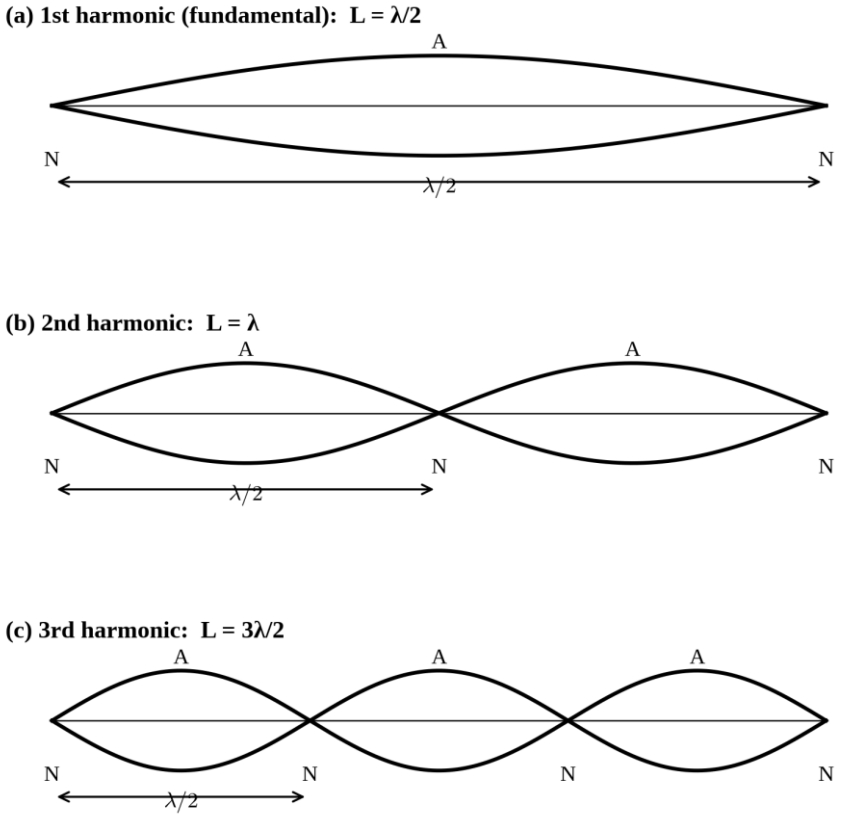


Figure: ***The first three harmonic patterns on a string fixed at both ends.*** *(a) 1st harmonic (fundamental): one half-wavelength fits between the two clamps, so L = λ/2. There are two nodes (one at each end) and one antinode (in the middle). (b) 2nd harmonic: two half-wavelengths fit, so L = λ. Three nodes, two antinodes. (c) 3rd harmonic: three half-wavelengths fit, so L = 3λ/2. Four nodes, three antinodes. In every harmonic the distance between adjacent nodes is λ/2.*

The wave speed on the string is fixed by its tension and linear mass density alone, so it is the same for every harmonic. That common wave speed times the wavelength gives the frequency:

$$f_n = \frac{v}{\lambda_n} = \frac{nv}{2L}, \quad n = 1, 2, 3, \ldots$$

Substituting $v = \sqrt{\frac{T}{\mu}}$ gives the master formula of the section:

$$\mathbf{f_n = \frac{n}{2L}\sqrt{\frac{T}{\mu}}, \quad n = 1, 2, 3, \ldots}$$

In particular, the fundamental frequency is:

$$f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$$

and every higher harmonic is an integer multiple of it: $\mathbf{f_n = nf_1}$. This integer-multiple structure is what makes a vibrating string a **harmonic** oscillator and not merely a periodic one. It is what allows a single bowed or plucked string to sound a definite musical pitch rather than an indefinite noise: the human ear hears the fundamental, while the ear and the brain together register the higher harmonics as the timbre that distinguishes a violin from a guitar from a *zeze*.

The three laws of string vibration

Read off the master formula three times, each time holding two of the three quantities (L, T, μ) fixed and varying the third.

Law 1 (length): *with T and μ fixed,* $f_1 \propto \frac{1}{L}$. A shorter string vibrates at a higher fundamental frequency. Halving the length doubles the fundamental: this is the rule a guitarist uses every time they press a string halfway along its length to play one octave higher.

Law 2 (tension): *with L and μ fixed,* $f_1 \propto \sqrt{T}$. A tighter string vibrates at a higher fundamental frequency, but only as the square root of the tension. Doubling the tension multiplies the fundamental by $\sqrt{2} \approx 1.41$, not by 2. This is why tuning a guitar requires fine adjustment: a small change in tension produces a small change in frequency.

Law 3 (linear mass density): *with L and T fixed,* $f_1 \propto \frac{1}{\sqrt{\mu}}$. A lighter string vibrates at a higher fundamental frequency. This is why the high strings of a guitar or piano are made of thin wire: lower μ gives higher frequency at the same length and tension.

All three laws together explain every stringed instrument ever made. To raise the pitch of a string at fixed material and fixed tension, shorten it (a violinist's left hand). To raise the pitch at fixed length and fixed material, tighten it (a tuning peg). To make a string play a low note without making it too long or too floppy, increase its mass per unit length (the wound copper of a piano bass string).

Adjacent resonances and the fundamental

A useful corollary of $f_n = nf_1$ concerns the spacing between consecutive harmonics. The difference between any two adjacent allowed frequencies on the same string is:

$$f_{n+1} - f_n = (n + 1)f_1 - nf_1 = f_1.$$

That is, **the difference between any two adjacent resonant frequencies of a fixed-fixed string is the fundamental frequency**. This corollary is exceedingly useful in the laboratory. If a sonometer string is observed to resonate at two adjacent driver frequencies, their numerical difference is the fundamental of the string, immediately, without any reference to length, tension, or mass density.

With the harmonic spectrum in hand and the three laws written down, let us put them through their paces using a few worked examples.

BINDER Example 23

A guitar string of length $L = 0.65\text{m}$ has total mass $m = 3.6\text{g}$ and is stretched to a tension $T = 480\text{N}$. Find: (a) its linear mass density μ, (b) the speed of transverse waves on the string, and (c) the fundamental frequency of the string.

Solution

(a) The linear mass density is mass per unit length:

$$\mu = \frac{m}{L} = \frac{3.6 \times 10^{-3}\ \text{kg}}{0.65\text{m}} = 5.5 \times 10^{-3}\ \text{kg/m}$$

(b) The speed of transverse waves follows from the string formula:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{480\ \text{N}}{5.5 \times 10^{-3}\ \text{kg/m}}} \approx 294\ \text{m/s}$$

(c) The fundamental frequency of a string fixed at both ends is:

$$f_1 = \frac{v}{2L} = \frac{294\ \text{m/s}}{2 \times 0.65\text{m}} \approx 226\ \text{Hz}$$

Making sense of the answer: *Three substitutions, three answers:* μ *from mass and length;* v *from* μ *and* T*;* f_1 *from* v *and* L*. The fundamental frequency 226Hz lies just above A3 (220Hz), the open A pitch on the cello and the twelfth-fret pitch of a guitar's A string. The number is squarely in the audible musical range, and it confirms that a string of mass 3.6g, length 65cm, and tension 480N produces a tone in the A region of the musical scale. Each of the three steps used exactly one and only one physical relation: the definition of linear mass density, the wave speed on a stretched string, and the fundamental frequency of a fixed-fixed string. There is no shortcut that skips a step; there is no step that needs more than one relation.*

Think like a physicist: *Notice the order in which the quantities appeared. The mass and length came together to give* μ*, a property of the string itself that does not depend on how the string is being used. The tension entered next, combining with* μ *to give* v*, a property of waves on this particular string at this particular tension. The length then re-entered (this time as a boundary), combining with* v *to select* f_1*, a property of the entire fixed-fixed system. The same string at the same tension would have a different fundamental if it were clamped at different points; the boundary condition is what turns a wave speed into a frequency. Hold this distinction. It is the difference between a property of the medium (which* v *is) and a property of the medium-plus-boundary system (which* f_1 *is).*

BINDER Example 24

A sonometer wire of length $L = 60\text{cm}$ and total mass $m = 12\text{g}$ is observed to carry a transverse wave of wavelength 4.0cm at a frequency of 720Hz. Find the tension in the wire.

Solution

Linear mass density:

$$\mu = \frac{m}{L} = \frac{0.012\ \text{kg}}{0.60\text{m}} = 0.020\text{kg/m}$$

Wave speed from $v = f\lambda$:

$$v = (720\ \text{Hz}) \times (0.040\ \text{m}) = 28.8\text{m/s}$$

Tension from $T = \mu v^2$:

$$T = (0.020\,\text{kg/m}) \times (28.8\,\text{m/s})^2 = 16.6\text{N}$$

Making sense of the answer: *The wave speed and tension followed from* $v = f\lambda$ *and* $T = \mu v^2$ *without any reference to which harmonic the wire happens to be vibrating in. The given wavelength 4.0cm fits* $L/(\lambda/2) = 0.60\text{m}/0.020\text{m} = 30$ *times into the wire, so the wire is in fact in its 30th harmonic, a high mode but a perfectly legitimate one. The tension calculation does not care about that, because tension and wave speed are properties of the wire itself, independent of which standing pattern is excited on it. A sonometer in this configuration is a routine laboratory setup: a tension of about 17N applied to a wire of mass density 20g/m supports waves at modest speed, and any of the discrete frequencies* $f_n = nf_1$ *can be excited by an external driver tuned to that frequency.*

Think like a physicist: *This is the* **inverse** *of the previous example. There, the unknowns were derived properties (wave speed, fundamental frequency) computed from the given physical setup (length, mass, tension). Here, one of those derived properties (the tension) is the unknown, and the input is a measurement of an actual standing pattern on the wire (its wavelength and frequency). The wave speed serves as the bridge:* $v = f\lambda$ *reads it off the observed pattern, and* $T = \mu v^2$ *converts it back to the underlying tension.*

BINDER Example 25

A steel piano wire of length $L = 0.50\text{m}$ has mass $m = 4.0\text{g}$ and is held under tension $T = 600\text{N}$. (a) Find the fundamental frequency of the wire. (b) Find the frequency of the third harmonic. (c) Find the highest harmonic of the wire whose frequency is at most 12000Hz.

Solution

Common preliminaries:

$$\mu = \frac{m}{L} = \frac{0.0040\text{ kg}}{0.50\text{m}} = 8.0 \times 10^{-3}\text{ kg/m},$$

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{600\text{ N}}{8.0 \times 10^{-3}\text{ kg/m}}} = 273.9\text{ m/s}.$$

(a) Fundamental frequency:

$$f_1 = \frac{v}{2L} = \frac{273.9\text{ m/s}}{2 \times 0.50\text{ m}} \approx 274\text{Hz}$$

(b) Third harmonic: $f_3 = 3f_1$:

$$f_3 = 3 \times 274\text{ Hz} = 822\text{Hz}$$

(c) The highest harmonic with frequency at most 12000Hz is the largest integer n satisfying,

$$nf_1 \leq 12000\text{Hz}$$

$$n \leq \frac{12000\text{ Hz}}{274\text{ Hz}} = 43.8, \quad \text{so} \quad n_{max} = 43$$

The highest harmonic is 43

Checking: The frequency of that harmonic is:

$f_{43} = 43 \times 274\text{Hz} = 11782\text{Hz}$ which is very close but less than the 12000Hz limit as required.

Making sense of the answer: *The wire's fundamental at about 274Hz lies near* $C_4\#$ *(277Hz), just above middle C. The third harmonic at about 822Hz is in the higher range a piano routinely sounds. The 43rd harmonic at 11.8kHz is a thin, high overtone, and audible to most young ears (the upper limit of human hearing is roughly 16-20kHz). A piano wire does not actually sound at any single one of these frequencies in isolation; struck, it sounds in a superposition of many of them, with the relative amplitudes determining the timbre that distinguishes a piano from a harpsichord and a harpsichord from a guitar.*

Think like a physicist: *Part (c) is a discrete counting problem dressed up as physics: divide a frequency limit by a fundamental, take the floor of the quotient, and that integer is the answer. Notice that* $n \times f_1$ *is a strictly increasing arithmetic sequence with common difference* f_1*, so once the fundamental is known the location of every harmonic on the frequency axis is fixed. The 43rd harmonic at 11.8kHz illustrates how rapidly a string can populate the upper end of the audible range: a wire whose fundamental is below middle C can produce overtones extending nearly to the edge of human hearing.*

BINDER Example 26

A string of length $L = 3.20\text{m}$, fixed at both ends, is observed to resonate at two adjacent driver frequencies, 170Hz and 204Hz, with no resonance for any driver frequency in between. Find: (a) the fundamental frequency of the string, and (b) the speed of transverse waves on the string.

Solution

(a) For a string fixed at both ends, the allowed frequencies are $f_n = nf_1$, and adjacent allowed frequencies differ by exactly the fundamental:

$$f_{n+1} - f_n = f_1$$

Therefore

$$f_1 = 204\text{Hz} - 170\text{Hz} = 34\text{Hz}$$

(b) The wave speed follows from $f_1 = v/(2L)$, solved for v:

$$v = 2L \times f_1 = 2 \times 3.20\text{m} \times 34\text{Hz} = 217.6\text{m/s}$$

Making sense of the answer: *The fundamental frequency and the wave speed both fell out of one number, the difference between two consecutive resonances, and the string length, with no reference to mass or tension. This is what makes the adjacent-resonance method useful in the laboratory: f_1 can be extracted from the resonance pattern alone, without weighing the string or measuring its tension. Once f_1 is known, the wave speed follows immediately from the length, and (if desired) the tension follows from $T = \mu v^2$ once the linear mass density is independently known. Note also which harmonic numbers correspond to the observed resonances: 170Hz = 5 × 34Hz and 204Hz = 6 × 34Hz, so the observed pair is the 5th and 6th harmonics of the string.*

Think like a physicist: *The three quantities L, T, and μ combine in the master formula to determine the fundamental frequency, but a single resonance frequency alone does not tell us which harmonic it is. Two adjacent resonances do: their difference is the fundamental, and the ratio of either to that fundamental is the harmonic number. This is the structural lesson. The harmonic spectrum is fixed by L, T, and μ; reading f_1 off the spectrum requires either knowing those three quantities directly, or sampling the spectrum at two adjacent points. The latter is often the easier route in practice.*

REAL Example 27

The physics laboratory at Miono Secondary, Friday afternoon. Mr. Akilikubwa walks in carrying two stringed instruments: a *kora*, the West African harp with twenty-one strings stretched between a long neck and a calabash gourd, and a *zeze*, a traditional Tanzanian one-string fiddle with a goat-skin resonator. He sets both on his desk and writes three relations on the board:

$$f_1 \propto \frac{1}{L}, \qquad f_1 \propto \sqrt{T}, \qquad f_1 \propto \frac{1}{\sqrt{\mu}}$$

Mr. Akilikubwa: *"Every stringed instrument that has ever moved a human being to tears (kora, zeze, violin, piano, guitar, harp, oud) was operating under exactly those three relationships. The one-string instrument I have here in my left hand and the twenty-one string instrument I have in my right hand differ in nothing else of physical importance. Their twenty-two strings, taken together, are twenty-two boundary-value problems. Each one obeys the same master formula. Each one selects a different fundamental by setting L, T, and μ to particular values that the maker chose deliberately."*

Kipute: *"Sir, when a guitarist presses a string halfway along its length and plucks it, what does the string do?"*

Kipanga (without prompting): *"It plays the second harmonic. The pressed-down finger forces a node at the midpoint, so the effective length of the freely vibrating segment is $L/2$. By the first law, f_1 of the shortened segment is twice the original f_1. The pitch jumps up one octave."*

Mr. Akilikubwa: *"He is learning. Notice what just happened, Kipute. Kipanga did not invoke the second harmonic of the original string. He invoked the first harmonic of a new, shorter string. The two are mathematically the same frequency in this particular case, but conceptually they are not the same thing. The original string's second harmonic is what you would hear if you plucked the original string and forced a node at the midpoint by holding it there gently with one finger while plucking with another. The shortened string's fundamental is what you hear when you press the finger down hard enough to clamp the string against the fret, terminating the original string and creating a new, shorter, fixed-fixed string. The numerical answer is the same; the physical setup is different."*

Kipute: *"And the kora's twenty-one strings, sir? They are all the same length on the instrument; the maker did not change L."*

Mr. Akilikubwa: *"Correct. The kora maker varies μ and T, not L. A traditional kora uses fishing-line strings of varying gauge: the bass strings are thicker (higher μ), the treble strings are thinner (lower μ), and each one is tuned to its desired fundamental by adjusting the leather tuning rings on the neck (varying T). The third law, $f_1 \propto 1/\sqrt{\mu}$, is what makes a low note possible on a string of fixed length without making the string impossibly long: increase μ instead. The piano maker does the same: the wound copper of a bass string raises μ enough that a low fundamental can be produced from a string short enough to fit inside the piano case."*

Use the dialogue to answer the following questions:

(a) The fundamental frequency of a particular kora bass string is 110Hz. The string is replaced by another string of the same length and tension but with twice the linear mass density. Find the new fundamental frequency.

(b) The high-pitched first string of a guitar plays f_1 = 330Hz when open. The guitarist presses the string against the seventh fret, which shortens the vibrating segment to $2/3$ of its original length. Assuming T and μ are unchanged, find the new fundamental frequency.

Solution

(a) By the third law, $f_1 \propto 1/\sqrt{\mu}$. Doubling μ multiplies f_1 by $1/\sqrt{2}$:

$$f_{1,\text{new}} = \frac{110\text{Hz}}{\sqrt{2}} = 77.8\text{Hz}$$

(b) By the first law, $f_1 \propto 1/L$. Reducing L to $2L/3$ multiplies f_1 by $3/2$:

$$f_{1,\text{new}} = \frac{3}{2} \times 330\text{Hz} = 495\text{Hz}$$

Making sense of the answer: *Both calculations used a single law each, applied to a single change of variable, with the other two variables held fixed. (a) Doubling the linear mass density lowers the fundamental by the factor $1/\sqrt{2}$, from 110Hz to about 78Hz: a kora bass string given a heavier replacement plays a lower note. (b) Pressing the high E string at the seventh fret shortens the vibrating segment to two-thirds of its open length, and the inverse-length law gives a fundamental of 495Hz: a perfect fifth above the open string's 330Hz. The factor 3:2 between the two frequencies is the famous **just intonation** interval that the ancient Pythagoreans first identified, and the same factor that every guitar's seventh fret produces. The maths and the music agree.*

Think like a physicist: *Mr. Akilikubwa's emphasis is the right one. The three laws each hold **only when the other two variables are fixed**. The kora maker who replaces a string with a heavier one of the same gauge but a different alloy is changing μ alone (Law 3). The guitarist who presses a fret is changing L alone (Law 1). The tuner who turns a peg is changing T alone (Law 2). Multiple simultaneous changes do not obey any one of the three simple proportionalities; they obey the master formula $f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$ and require the master formula's full machinery. The three laws are the partial derivatives of that master formula in disguise; treating them as more than that is a recipe for arithmetic that does not match what the instrument actually does.*

HOT Example 28

A stretched string clamped at both ends is sounded together with a tuning fork while the string is held under tension $T_1 = 245\text{N}$. The combination produces beats at 15 beats per second. The tension on the string is then increased smoothly to $T_2 = 320\text{N}$, and at that tension the string and the fork are heard to be in unison: no beats. Find the frequency of the tuning fork.

Solution

Let f_1 denote the string's fundamental frequency at T_1, f_2 its fundamental frequency at T_2, and F the (constant) frequency of the fork. By the second law of string vibration, with L and μ fixed, $f_1 \propto \sqrt{T}$, so:

$$\frac{f_2}{f_1} = \sqrt{\frac{T_2}{T_1}} = \sqrt{\frac{320\text{N}}{245\text{N}}} = \sqrt{\frac{64}{49}} = \frac{8}{7}$$

At T_2 the string is in unison with the fork, so $f_2 = F$.
At T_1 the beat condition is $|f_1 - F| = 15\text{Hz}$.
The string's frequency rose with the tension to meet the fork's frequency, so f_1 was below F, and we drop the absolute value:

$$F - f_1 = 15\text{Hz}$$

Substitute $F = f_2 = \frac{8}{7}f_1$:

$$\frac{8}{7}f_1 - f_1 = 15\text{ Hz}, \quad \text{so} \quad f_1 = 105\text{ Hz}.$$

Therefore:

$$F = \frac{8}{7} \times 105\text{Hz} = 120\text{Hz}$$

The frequency of the tuning fork is 120Hz.

Making sense of the answer: *The fork frequency is 120Hz, near the lower end of the bass-baritone vocal range. The string's fundamental rose from 105Hz at T_1 to 120Hz at T_2, a 15Hz increase that matches the original 15Hz beat frequency, as it must. The arithmetic worked out cleanly because the two tensions were in the ratio $320:245 = 64:49$, whose square root is the rational number $8:7$. Real-world tunings rarely give such clean ratios; we have chosen them here so that the structure of the calculation, and not the arithmetic, is the lesson.*

Think like a physicist: *The crucial physical step in this example is not the use of $f_1 \propto \sqrt{T}$, that is straightforward. The crucial step is deciding which side of the unison the string was on at T_1. The beat frequency 15Hz tells us how far apart the string and the fork were in absolute value, but not which one was higher. The physical clue is the direction of the tension change. Tension was **increased** between T_1 and T_2, which **raised** the string's frequency. If raising the string's frequency brought it into unison with the fork, the fork must have been **higher** than the string at the start. That single sentence resolved the ambiguity in the absolute value and let us drop it. Always look for the physical clue that selects between the algebraic possibilities; algebra alone will not do it.*

HOT Example 29

A monochord (a single-string instrument with a movable bridge for changing the vibrating length) is to be built so that its open-string fundamental is $f_1 = 220\text{Hz}$. The wire available for the string has linear mass density $\mu = 2.5 \times 10^{-3}\ \text{kg/m}$ and a maximum safe operating tension $T_{max} = 450\text{N}$, above which the wire is at risk of snapping.

(a) At a tension $T = 300\text{N}$, find the speed of transverse waves on the wire.
(b) For $T = 300\text{N}$, find the string length L that gives the required fundamental frequency.
(c) For a vibrating length $L = 40\text{cm}$, find the tension that gives the required fundamental frequency.
(d) Is the choice $L = 40\text{cm}$ within the safe operating range of the wire?
(e) What is the longest vibrating length that gives the required fundamental, given the maximum safe tension T_{max}?

Solution

(a) Wave speed at $T = 300\text{N}$:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{300\ \text{N}}{2.5 \times 10^{-3}\ \text{kg/m}}} \approx 346\ \text{m/s}$$

(b) Length for $f_1 = 220\text{Hz}$ at $v \approx 346\ \text{m/s}$:

$$L = \frac{v}{2f_1} = \frac{346\ \text{m/s}}{2 \times 220\ \text{Hz}} = 0.787\text{m}, \quad \text{(about 78.7cm)}$$

(c) Tension for $L = 0.40\ \text{m}$ at $f_1 = 220\ \text{Hz}$. First the wave speed:

$$v = 2L \times f_1 = 2 \times 0.40\text{m} \times 220\text{Hz} = 176\text{m/s}$$

Then the tension from $T = \mu v^2$:

$$T = (2.5 \times 10^{-3}\ \text{kg/m}) \times (176\ \text{m/s})^2 = 77.4\ \text{N}$$

(d) Compare the required tension with the maximum:

$$77.4\ \text{N} < T_{max}\ (450\ \text{N})$$

Hence the choice is well within the safe operating range.
The tension required by $L = 40\ \text{cm}$ uses only about 17% of the wire's strength.

(e) Maximum wave speed at the maximum tension:

$$v_{max} = \sqrt{\frac{T_{max}}{\mu}} = \sqrt{\frac{450\text{N}}{2.5 \times 10^{-3}\ \text{kg/m}}} \approx 424\ \text{m/s}.$$

Maximum vibrating length at T_{max} for the same fundamental:

$$L_{max} = \frac{v_{max}}{2f_1} = \frac{424\ \text{m/s}}{2 \times 220\ \text{Hz}} = 0.964\text{m}, \quad \text{(about 96.4cm)}$$

Making sense of the answer:

The design space is bounded. To play a 220Hz fundamental on this wire, the vibrating length must be between zero and about 96cm; below 96cm the required tension is below the wire's safe limit, and above 96cm the required tension exceeds it and the wire is at risk of snapping. The numerical answers show how design freedom is exercised: a short string (40cm) needs a low tension (77N); a medium string (79cm) needs a medium tension (300N); the longest possible string (96cm) needs the maximum safe tension (450N). The designer chooses based on instrument size, the desired feel of the strings under the player's fingers, and the safety margin.

Think like a physicist: *The fundamental frequency depends on three quantities (L, T, μ) combined into one master formula. With μ fixed by the choice of wire and f_1 fixed by the design specification, the designer has two free parameters left, L and T, and one constraint. That leaves a one-dimensional family of solutions: every (L, T) pair on a particular curve in the (L, T) plane is a valid design. The maximum-tension constraint truncates that curve at one end; the minimum-tension constraint (the wire must still be taut) would truncate it at the other. The designer picks any point on the surviving curve and builds an instrument. This is how engineering design actually proceeds: the physics gives the family of possibilities; the engineer picks one based on practical considerations the physics does not encode.*

The harmonic spectrum of a fixed-fixed string is a discrete sequence $f_n = nf_1$, with $f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$, and every higher harmonic an integer multiple of the fundamental. Three laws follow by holding two of the three quantities fixed and varying the third. Every stringed musical instrument and every sonometer in every laboratory operates within those laws. The next section generalises the wave-speed formula beyond strings to all elastic media, and asks what limits the speed of a transverse wave on a real wire of finite tensile strength.

VELOCITY OF MECHANICAL WAVES IN DIFFERENT MEDIA

The previous section gave a single formula, $v = \sqrt{\frac{T}{\mu}}$, for the speed of a transverse wave on a stretched string. The formula is the launching point for a far broader statement that holds for every elastic medium in classical mechanics. The structure of the formula is what generalises: *a square root, with an **elastic property** of the medium in the **numerator** and an **inertial property** of the medium in the **denominator***. Strings, solids, and liquids all obey it; only the specific elastic and inertial properties change from medium to medium. This section presents the generalised formula, applies it to longitudinal waves in solids and liquids, and asks what physical limit caps the transverse wave speed on a real wire of finite tensile strength.

Gases require one further refinement (an adiabatic factor that Newton missed and Laplace later supplied), and they are postponed to Chapter 9 where sound in air will be treated in full. This section confines itself to the media where the elastic property is straightforward: **tension** on a *string*, **Young's modulus** in a *solid*, and **bulk modulus** in a *liquid*.

The unified wave-speed formula

In every elastic medium that supports a mechanical wave, the speed of that wave is the square root of an elastic property divided by an inertial property of the medium:

$$\mathbf{v} = \sqrt{\frac{\textbf{elastic property}}{\textbf{inertial property}}}$$

The structure is a piece of physical bookkeeping. The **numerator** *quantifies how stiffly the medium resists deformation:* the **larger the elastic property**, the larger the restoring force the medium exerts when displaced from equilibrium, and the **faster a disturbance propagates**. The **denominator** *quantifies how heavy the medium is per unit volume (or per unit length, in the case of a string):* the **larger the inertia**, the more sluggish the medium's response to any restoring force, and the **slower the wave propagates**. The square root captures the way these two competing tendencies enter the equation of motion: every elastic medium obeys a wave equation in which the wave speed enters as the square root of the ratio.

Three particular cases cover the bulk of this chapter and the next:

1) Transverse waves on a stretched string

Elastic property: the *tension T* (force restoring the string toward its equilibrium line). Inertial property: the *linear mass density* μ (mass per unit length). The wave speed is:

$$v = \sqrt{\frac{T}{\mu}}$$

2) Longitudinal waves in a solid rod

Elastic property: *Young's modulus Y* (force per unit area per unit fractional extension along the rod). Inertial property: the *volume mass density* ρ (mass per unit volume). The wave speed is:

$$v = \sqrt{\frac{Y}{\rho}}$$

3) Longitudinal waves in a liquid

Elastic property: the *bulk modulus K* (pressure required to compress the liquid by unit fractional volume). Inertial property: the *volume mass density* ρ. The wave speed is:

$$v = \sqrt{\frac{K}{\rho}}$$

All three are special cases of the same square-root rule. Once a student has internalised the rule, the only thing left to learn for any new medium is which elastic property and which inertial property apply: from there the wave speed follows by one substitution.

The maximum transverse wave speed on a real wire

A real stretched wire cannot be tightened indefinitely. Beyond a certain tensile stress (force per unit cross-sectional area), the wire's atomic bonds yield and the wire snaps. The maximum stress a wire can sustain without breaking is its **ultimate tensile stress**, denoted $\boldsymbol{\sigma_{max}}$. This finite limit places a finite upper bound on the wave speed that the wire can carry, and the bound has a particularly elegant form.

Take a uniform wire of cross-sectional area A, linear mass density μ, and volume mass density ρ. The two densities are connected by $\mu = \rho A$: a length ℓ of wire has volume $A\ell$ and mass $\rho A\ell$, so its mass per unit length is ρA. The maximum tension is the maximum stress times the area: $T_{max} = \sigma_{max} A$. Substituting both into $v = \sqrt{\frac{T}{\mu}}$:

$$\mathbf{v_{max}} = \sqrt{\frac{T_{max}}{\mu}} = \sqrt{\frac{\sigma_{max}A}{\rho A}} = \sqrt{\frac{\boldsymbol{\sigma_{max}}}{\boldsymbol{\rho}}}$$

The cross-sectional area A cancels exactly. The maximum transverse wave speed depends only on the wire's material, through σ_{max} and ρ, and not on its thickness. A thicker wire of the same material can take more tension and is also heavier per unit length; the two effects scale identically with A and cancel. This cancellation is the elegant structural fact that closes the section: *for any given material there is a single maximum wave speed, set by its ultimate tensile stress and its density, and that speed cannot be exceeded by any wire made of that material no matter what gauge it is.*

With the unified formula now stated for any elastic medium, the worked examples can begin.

BINDER Example 30

A long uniform steel rod is struck on one end with a hammer. Find the speed of the longitudinal disturbance that propagates along the rod, given that steel has Young's modulus $Y = 2.0 \times 10^{11}\text{Pa}$ and density $\rho = 7800\text{kg/m}^3$.

Solution

For a longitudinal wave in a solid rod, the elastic property is Young's modulus and the inertial property is the volume mass density:

$$v = \sqrt{\frac{Y}{\rho}} = \sqrt{\frac{2.0 \times 10^{11}\text{Pa}}{7800\text{kg/m}^3}} = 5.06 \times 10^3\text{m/s}$$

That is approximately 5060m/s, or about 15 times the speed of sound in air at room temperature (340m/s).

Making sense of the answer: *The arithmetic is one substitution into $v = \sqrt{\frac{Y}{\rho}}$. The result, about 5060m/s, is in close agreement with measured values for the speed of longitudinal sound in steel rods, which experiment places between 5000m/s and 5200m/s depending on the specific alloy. The reason steel structures ring so loudly when struck (a steel rail clanged with a hammer, the side of a steel ship, the bell of a railway crossing) is that the longitudinal wave travels through the steel with this very high speed and re-radiates as audible air-pressure oscillations at the surface of the steel. The sharpness of the ring follows from the speed: a struck steel object delivers its acoustic energy to the air in a brief pulse, not in the slow diffusion characteristic of less stiff materials.*

Think like a physicist: *Compare this calculation with the string-tension calculation of the previous section. The structure is identical: square root, elastic property over inertial property, one substitution. Only the names of the properties have changed. For the string the elastic property was a force (T); for the solid it is a force per unit area per unit strain (Y). For the string the inertial property was a mass per unit length (μ); for the solid it is a mass per unit volume (ρ). The dimensions of the ratio still work out to those of speed squared, and the square root recovers the speed. Whenever a new elastic medium appears, the question to ask is which elastic property restores it and which inertial property loads it. From those two answers the wave speed follows automatically.*

REAL Example 31

On the Tanzania-Zambia (TAZARA) railway line, before the era of two-way radios and satellite tracking, track-walkers and trackside workers used a centuries-old technique to detect an approaching train long before they could hear it through the air: they pressed one ear to the steel rail and listened. The vibration carried by the steel arrived noticeably earlier than the sound carried by the air, giving the listener several seconds of warning to step clear. Given that:

Steel: $Y = 2.0 \times 10^{11}\text{Pa}$, $\rho = 7800\text{kg/m}^3$. Speed of sound in air: $v_{air} = 340\text{m/s}, K_{air} = 1.4 \times 10^5\text{Pa}, \rho_{air} = 1.2\text{kg/m}^3$

(a) Calculate the speed of the longitudinal vibration that travels through the steel rail.
(b) A train is 1.5km up the line from a worker pressing his ear to the rail. How much earlier does the vibration through the steel reach the worker than the sound through the air?
(c) Explain physically, using the unified wave-speed formula, why sound travels so much faster in steel than in air.

Solution

(a) From the previous example, $v_{steel} = \sqrt{\frac{Y}{\rho}} \approx 5060\text{m/s}$.

(b) Distance $d = 1500\text{m}$. Time through the steel:

$$t_{steel} = \frac{d}{v_{steel}} = \frac{1500\text{m}}{5060\text{m/s}} = 0.296\text{s}$$

Time through the air:

$$t_{air} = \frac{d}{v_{air}} = \frac{1500\text{m}}{340\text{m/s}} = 4.41\text{s}$$

The advance warning is the difference:

$$\Delta t = t_{air} - t_{steel} \approx 4.41\text{s} - 0.30\text{s} = 4.1\text{ seconds.}$$

(c) The unified formula is $v = \sqrt{\frac{\text{elastic}}{\text{inertia}}}$. Steel's Young's modulus is approximately 2×10^{11}Pa; the bulk modulus of air at atmospheric pressure is approximately 1.4×10^{5}Pa. Steel is therefore about 1.4×10^{6} times stiffer than air as an elastic medium. Steel's density is 7800kg/m^3; air's density at room temperature is about 1.2kg/m^3. Steel is therefore about 6500 times denser than air. The ratio inside the square root, (elastic/inertia), is therefore about $\frac{1.4 \times 10^6}{6500} \approx 220$ times larger for steel than for air. The wave speed scales with the square root of that ratio, so the speed in steel is about $\sqrt{220} \approx 15$ times the speed in air. Steel's overwhelming stiffness advantage easily defeats its density disadvantage, and hence the sound waves runs faster in steel.

Making sense of the answer: *The 4-second warning is real and was decisive in the days before mechanical track sensors. A worker bent over the rail with his ear pressed to the steel learned of the approaching train roughly 4.1 seconds before the same train would have been audible through the air at the same distance, and that interval was the difference between safety and disaster. The ratio of speeds, 15:1, is the structural answer: steel is so much stiffer than air that even though it is also much denser, the elastic-over-inertial ratio comes out far in favour of steel.*

Think like a physicist: *This is one of the few cases in physics where two factors point in opposite directions but one factor wins by a wide enough margin that the qualitative answer is obvious. Steel is stiffer than air; this raises the wave speed. Steel is denser than air; this lowers the wave speed. The two effects compete inside the square root. In some materials the competition is genuinely close, and the answer requires a careful numerical calculation. Always check both factors when comparing wave speeds in two media.*

HOT Example 32

A wire is made of a metal alloy with ultimate tensile stress $\sigma_{max} = 5.0 \times 10^{8}$Pa and density $\rho = 8000\text{kg/m}^3$. The wire has a cross-sectional area $A = 2.0 \times 10^{-6}\text{m}^2$.

(a) Find the maximum transverse wave speed the wire can carry without snapping, using the formula $v_{max} = \sqrt{\frac{\sigma_{max}}{\rho}}$.

(b) Find the maximum tension and the linear mass density of the wire.

(c) Compute the wave speed from $v = \sqrt{\frac{T}{\mu}}$ using the values from part (b), and verify that the answer agrees with part (a).

Solution

(a) Maximum transverse wave speed:

$$v_{max} = \sqrt{\frac{\sigma_{max}}{\rho}} = \sqrt{\frac{5.0 \times 10^{8}\text{Pa}}{8000\text{kg/m}^3}} = 250\text{m/s}$$

(b) Maximum tension and linear mass density:

From the definition of tensile stress, $T_{max} = \sigma_{max}A$:

$$T_{max} = (5.0 \times 10^{8}\text{Pa}) \times (2.0 \times 10^{-6}\text{m}^2) = 1000\text{N}$$

From $\mu = \rho A$:

$$\mu = (8000\text{kg/m}^3) \times (2.0 \times 10^{-6}\text{m}^2) = 0.016\text{kg/m}$$

(c) Verify by direct substitution into the string formula:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{1000\text{N}}{0.016\text{kg/m}}} = 250\text{m/s}$$

The two computations agree exactly, as the algebra of the section guarantees they must.

Making sense of the answer: *The two computations agree because the cancellation of A between numerator and denominator is exact, not approximate. Algebraically, $\frac{\sigma_{max}A}{\rho A} = \frac{\sigma_{max}}{\rho}$, so the cross-sectional area drops out completely and the maximum wave speed is a property of the material alone. A thicker wire of the same alloy has the same maximum transverse wave speed as a thinner wire: the thicker wire can take more tension (because $T_{max} \propto A$), but it is also heavier per unit length (because $\mu \propto A$), and the two effects scale together so as to leave the wave speed unchanged.*

Think like a physicist: *This kind of cancellation is not a coincidence. Whenever a physical quantity is the ratio of two extensive properties (properties that scale with the size of the system), the scaling factors cancel and the quantity becomes intensive (a property of the material alone, independent of system size). Tension is extensive; mass is extensive; their ratio under the square root, divided by the cross-sectional area in numerator and denominator, leaves an intensive ratio of stress to density.*

The unified wave-speed formula $v = \sqrt{\frac{\text{elastic property}}{\text{inertial property}}}$ governs every mechanical wave in every elastic medium, with strings, solids, and liquids each providing one specialisation. Real wires of finite tensile strength have a finite maximum transverse wave speed determined by the material alone, $v_{max} = \sqrt{\frac{\sigma_{max}}{\rho}}$, a result whose elegance comes from the exact cancellation of the cross-sectional area. Gases supply one further refinement that Newton missed and Laplace corrected, and they belong to Chapter 9 on sound.

The next section brings together the results we have developed and shows how they apply to real-world situations and everyday life.

APPLICATIONS OF WAVE MOTION

So far, the chapter has focused on building the theory of waves through definitions, derivations, and worked examples. This section now turns from the page to the world beyond it. It asks what these results actually do once they leave the classroom and enter real life. The purpose of the final section is therefore not to introduce new physics, but to reveal how remarkably far the physics already in hand can reach.

1. Musical instruments

Every stringed instrument that has ever been built (the kora, the zeze, the violin, the guitar, the piano, the harp, the oud) is a collection of fixed-fixed strings tuned by the maker to particular fundamentals. The harmonic spectrum $f_n = nf_1$, with $f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$, determines the pitch the string sounds. The three laws (length, tension, linear mass density) determine how the maker varies the strings to produce a chosen range of pitches: a piano lengthens the bass strings and weights them with copper winding, a violin shortens its strings and uses higher tension, a guitar uses six strings of varying mass density tuned to fixed pitches. None of this is discretionary; it is fixed by the master formula and the boundary condition that both ends of a string under a player's finger are nodes.

Two further pieces complete the picture. First, when a string is plucked or bowed it does not vibrate at any single one of the allowed harmonics; it vibrates in a superposition of many of them, with the relative amplitudes set by where and how it is excited. The mixture of harmonics, called the **timbre**, is what tells a listener at the back of a concert hall that a violin is playing rather than a guitar even when the two are sounding the same fundamental. Second, the same boundary-condition selection that picks out the harmonics on a string also picks out the resonant modes of an air column inside an organ pipe, an open or closed flute, or a clarinet. The full treatment of pipe modes belongs to Chapter 9; the pattern (a discrete spectrum of allowed frequencies, set by boundary conditions on the standing wave) is identical to the one this chapter has worked out for strings.

2. Seismology and the structure of the Earth

When an earthquake releases its energy near the surface of the Earth, the disturbance travels outward through the planet as a mechanical wave. Two distinct kinds of wave propagate. The **primary** wave (P-wave) is longitudinal: the rock oscillates back and forth along the direction of propagation, the way a sound wave does in air. The **secondary** wave (S-wave) is transverse: the rock oscillates perpendicular to the direction of propagation, the way a wave does on a guitar string. Each obeys the unified speed formula $v = \sqrt{\frac{\text{elastic}}{\text{inertia}}}$, but with different elastic moduli in the numerator. The P-wave's restoring force comes from a combination of the rock's bulk modulus and its shear modulus; the S-wave's restoring force comes from the shear modulus alone. The mathematical consequence is that P-waves always travel faster than S-waves in the same rock.

Two consequences follow that have shaped the science of geophysics. First, when an earthquake's waves arrive at a distant seismograph, the P-wave arrives first and the S-wave arrives later, and the gap between their arrival times reads back the distance to the epicentre: the longer the gap, the further the source. With three or more seismographs at known locations, the position of the epicentre is fixed by triangulation. Second, S-waves cannot propagate through a liquid, because a liquid has no shear strength. The fact that S-waves are absent from seismograph records on the side of the Earth opposite a major earthquake is what told twentieth-century geophysicists that the Earth has a liquid outer core. The diagnosis was made entirely from the absence of a kind of wave whose physics this chapter has already worked out.

3. Medical ultrasound imaging

A medical ultrasound scanner sends short pulses of progressive longitudinal waves into the body, at frequencies between roughly 1MHz and 20MHz, and listens for the echoes. Each tissue boundary (skin to fat, fat to muscle, muscle to bone, the wall of an organ, the wall of a foetal heart) reflects part of the wave back toward the transducer. The pulse-echo timing converts directly to depth: a wave that returns Δt seconds after it was emitted has travelled $v\Delta t$ through tissue and back, so the reflecting boundary is at depth $d = v\Delta t/2$, with v the speed of sound in soft tissue (about 1540m/s, set by the unified formula with the bulk modulus and density of water-rich biological tissue).

Two design choices are dictated by the physics of the chapter. The first is the choice of frequency. Higher frequency means shorter wavelength means finer spatial resolution; but higher frequency also means greater absorption in the tissue and shorter penetration depth. Obstetric scans of a foetus several centimetres deep use frequencies near the lower end of the range; superficial scans of a thyroid or a tendon a few millimetres deep use frequencies near the upper end. The second is the duration of the pulse: short pulses give better depth resolution, but a pulse can be no shorter than a few wave periods, so the depth resolution is fundamentally limited by the wavelength. Both trade-offs are direct consequences of $v = f\lambda$ applied to a medium whose absorption increases with frequency.

4. Resonance in bridges and buildings

Every elastic structure (a footbridge, a road bridge, a tall building, a chimney, a guitar's wooden body) has a discrete spectrum of natural frequencies of vibration, set by the same kind of boundary-condition selection that gives a string its harmonics. The structure's lowest natural frequency is called its fundamental; the higher natural frequencies are its overtones. If the structure is driven by an external force whose frequency matches one of these natural frequencies, the amplitude of the structure's response builds up dramatically through resonance, and the resulting oscillation can be large enough to damage or destroy the structure.

Foot traffic across a footbridge can drive the bridge's lateral mode if the footfalls happen to land at the natural frequency, with one well-known modern example being the synchronised lateral sway of a London footbridge on its opening day in the year 2000. Marching regiments are routinely instructed to break step when crossing a small bridge for exactly this reason. Wind blowing across a tall building can drive its bending modes through vortex shedding at frequencies determined by the wind speed and the building's cross-section. Modern civil engineering specifies a structure's natural frequencies during design, identifies the loading conditions that could excite each mode, and adds damping or stiffening so that no excitable mode lies near a frequency the building or bridge is likely to encounter in service. The physics is the same boundary-value problem that selects the harmonics of a string, scaled up from a few millimetres of wire to several hundred metres of steel.

5. Active noise cancellation

Noise-cancelling headphones, the cabin of a modern aircraft, and certain car engine compartments use a direct application of the principle of superposition. A microphone outside the listener's ear samples the unwanted sound. A small electronic processor inverts the sampled waveform (shifting its phase by π) and a tiny loudspeaker in the ear cup re-emits the inverted wave in real time. The original wave and the inverted wave superpose at the listener's eardrum; their displacements sum to zero, exactly the destructive interference we worked out in the chapter on superposition. The listener hears silence (or near silence) where there would otherwise have been engine rumble.

The technique works best for low-frequency steady noise: the inversion-and-replay loop has to complete in less than a fraction of a wavelength, and shorter wavelengths leave the loop too little time to catch up with the incoming wave. High-pitched sudden sounds, like a baby's cry or a screech of brakes, are not well cancelled and remain audible. The physics is unmistakable: superposition gave us **constructive interference** when two waves added with the same phase and **destructive interference** when they added with opposite phase. The headphones simply manufacture the second condition deliberately.

Having tasted each section on its own, we are now ready for the combination plate: miscellaneous worked examples that bring the whole chapter together on a single tray.

MISCELLANEOUS WORKED EXAMPLES ON WAVE MOTION

Example 33

(a) Account for the observation that water surface waves are classified as transverse, while sound waves travelling through the same water are longitudinal.

(b) A long horizontal spring is laid on a smooth bench and one end is moved back-and-forth along the spring's own length, sending a train of longitudinal waves along it. The compressions in the wave train are observed to be 0.40m apart, and a particular compression is timed to travel 6.0m along the spring in 0.50s. Find the wave speed, the frequency, the period, the angular frequency and the wave number of the wave.

Solution

(a) The classification depends on the direction in which the molecules of the medium are forced to oscillate, not on the medium itself. A water surface wave is held together by gravity acting on the displaced surface; the surface has to rise and fall for the restoring force to act, so each surface molecule oscillates vertically while the disturbance propagates horizontally, so the disturbance is perpendicular to the propagation and the wave is therefore transverse. A sound wave in the same body of water is held together by the bulk elasticity of water; compressing the water momentarily generates a restoring force that drives the molecules along the line of compression, so each molecule oscillates along the direction of propagation, so the disturbance is parallel to the propagation and the wave is therefore longitudinal. The two waves coexist in the same water without confusion because they exploit two independent restoring mechanisms.

(b) Wavelength is the distance between consecutive compressions, $\lambda = 0.40\text{m}$.

Wave speed:

$$v = \frac{6.0\text{m}}{0.50\text{s}} = 12\text{m/s}$$

Frequency:

$$f = \frac{v}{\lambda} = \frac{12\text{m/s}}{0.40\text{m}} = 30\text{Hz}$$

Period:

$$T = \frac{1}{f} = \frac{1}{30\text{Hz}} = 0.033\text{s}$$

Angular frequency:

$$\omega = 2\pi f = 2\pi(30\text{Hz}) = 60\pi\text{rad/s} \approx 188\text{rad/s}$$

Wave number:

$$k = \frac{2\pi}{\lambda} = \frac{2\pi}{0.40\text{m}} = 5\pi\text{rad/m} \approx 15.7\text{rad/m}$$

Example 34

(a) A leaf placed on the surface of a still pond bobs up and down without travelling along with the ripples that pass beneath it. Account for this in physical terms.

(b) A wave is described by $y(x,t) = 0.005\sin(2\pi x - 100\pi t)\text{m}$, with x in metres and t in seconds. State the direction of propagation and find the wavelength, the period, the wave speed and the maximum particle speed of the wave.

Solution

(a) A wave is a propagating disturbance, not a flow of matter. The water beneath the leaf executes a small simple-harmonic up-and-down motion as each ripple passes; the molecules return to where they started by the time the ripple has gone. The leaf, riding on the water, shares only that local up-and-down motion. The energy and the geometric pattern of the disturbance travel outward, but the water (along with the leaf) remains on average in place.

(b) The argument $2\pi x - 100\pi t = -(100\pi t - 2\pi x)$, so $y = -0.005\sin(100\pi t - 2\pi x)\text{m}$. Comparing with $y = A\sin(\omega t - kx)$, the wave travels in the $+x$ direction with $\omega = 100\pi\text{rad/s}$ and $k = 2\pi\text{rad/m}$.

Wavelength:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{2\pi\text{rad/m}} = 1.0\text{m}$$

Period:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{100\pi\text{rad/s}} = 0.020\text{s}$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{100\pi\text{rad/s}}{2\pi\text{rad/m}} = 50\text{m/s}$$

Maximum particle speed:

$$v_{p,max} = A\omega = (0.005\text{m})(100\pi\text{rad/s}) = 0.5\pi\text{m/s} \approx 1.57\text{m/s}$$

Example 35

(a) A diver swimming a few metres beneath the surface of a calm lake hears almost nothing of the loud conversation between two friends standing in a boat directly above her. Account for this observation.

(b) A sound wave of frequency 1.50kHz travels through fresh water at 1500m/s. It then crosses a clear boundary into a thick layer of mud at the bottom of the lake, where its speed is 1100m/s. Find the wavelength of the wave in the water, the wavelength in the mud and the percentage change in wavelength.

Solution

(a) Sound moving from air into water crosses a boundary across which the elastic and inertial properties of the medium change abruptly: the product of density and wave speed (the acoustic impedance) is roughly four thousand times larger in water than in air. A boundary with such a large impedance mismatch reflects almost all the incident energy back into the originating medium and transmits only about 0.1% of it across. The diver therefore receives a tiny fraction of the air-borne sound, which is why she hears almost nothing.

(b) The source frequency is preserved across the boundary; the medium changes the speed and the wavelength adjusts accordingly.

Wavelength in water:

$$\lambda_1 = \frac{v_1}{f} = \frac{1500\text{m/s}}{1500\text{Hz}} = 1.00\text{m}$$

Wavelength in mud:

$$\lambda_2 = \frac{v_2}{f} = \frac{1100\text{m/s}}{1500\text{Hz}} = 0.733\text{m}$$

Percentage change:

$$\frac{\lambda_2 - \lambda_1}{\lambda_1} \times 100\% = \frac{0.733\text{m} - 1.00\text{m}}{1.00\text{m}} \times 100\% = -26.7\%$$

The wavelength decreases by about 26.7% on crossing the boundary into mud.

Example 36

(a) Account for the experimental observation that the colour of a beam of light passing from air into a glass of water does not change, even though both its speed and its wavelength decrease.

(b) Light of wavelength 600nm in vacuum enters a slab of glass of refractive index 1.50. Find the wavelength and the frequency of the light inside the glass. Take $c = 3.00 \times 10^8\text{m/s}$.

Solution

(a) The colour seen by the eye is determined by the frequency of the light, not by its wavelength. Frequency is the rate at which the source emits oscillations of the electromagnetic field; conservation of oscillation count requires the same rate to emerge from any boundary the wave crosses. The medium can change the speed of the wave, and the wavelength is forced to scale in step ($\lambda = v/f$), but the frequency, and hence the colour, is preserved.

(b) Frequency in vacuum (preserved on entering the glass):

$$f = \frac{c}{\lambda_0} = \frac{3.00 \times 10^8\text{m/s}}{600 \times 10^{-9}\text{m}} = 5.00 \times 10^{14}\text{Hz}$$

Speed of light in the glass:

$$v_{\text{glass}} = \frac{c}{n} = \frac{3.00 \times 10^8\text{m/s}}{1.50} = 2.00 \times 10^8\text{m/s}$$

Wavelength inside the glass:

$$\lambda_{\text{glass}} = \frac{v_{\text{glass}}}{f} = \frac{2.00 \times 10^8\text{m/s}}{5.00 \times 10^{14}\text{Hz}} = 4.00 \times 10^{-7}\text{m} = 400\text{nm}$$

The frequency is unchanged at $5.00 \times 10^{14}\text{Hz}$ and the wavelength inside the glass is 400nm.

Example 37

(a) Distinguish between the wave velocity and the particle velocity of a transverse mechanical wave, indicating which depends on the amplitude of the wave and which does not.

(b) A transverse wave on a horizontal nylon rope is described by $y(x,t) = 0.030\sin(120\pi t - 4\pi x)\text{m}$, with x in metres and t in seconds. Find the wavelength, the wave speed, the displacement of the particle at $x = 0.50\text{m}$ at the instant $t = \frac{1}{240}\text{s}$, and the velocity of that same particle at the same instant.

Solution

(a) The wave velocity $v = \omega/k$ is the speed at which the wave pattern (each crest, each trough) advances through the medium. It is fixed by the elastic and inertial properties of the medium and does not depend on the amplitude. The particle velocity $v_p = \partial y/\partial t$ is the instantaneous speed at which a chosen particle of the medium oscillates about its own equilibrium position; its maximum value $v_{p,\max} = A\omega$ is directly proportional to the amplitude. Doubling the amplitude doubles the maximum particle speed but leaves the wave speed unchanged.

(b) Comparing $y = 0.030\sin(120\pi t - 4\pi x)\text{m}$ with $y = A\sin(\omega t - kx)$:

$$A = 0.030\text{m}, \qquad \omega = 120\pi\text{rad/s}, \qquad k = 4\pi\text{rad/m}$$

Wavelength:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{4\pi\text{rad/m}} = 0.50\text{m}$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{120\pi\text{rad/s}}{4\pi\text{rad/m}} = 30\text{m/s}$$

Phase at $x = 0.50\text{m}$, $t = \frac{1}{240}\text{s}$:

$$\omega t - kx = (120\pi)\left(\frac{1}{240}\right) - (4\pi)(0.50) = \frac{\pi}{2} - 2\pi = -\frac{3\pi}{2}$$

Displacement:

$$y = (0.030\text{m})\sin\left(-\frac{3\pi}{2}\right) = (0.030\text{m})(+1) = 0.030\text{m}$$

The particle is at the upper extreme of its oscillation.

Particle velocity:

$$v_p = A\omega\cos(\omega t - kx) = (0.030\text{m})(120\pi\text{rad/s})\cos\left(-\frac{3\pi}{2}\right) = (3.6\pi\text{m/s})(0) = 0$$

The particle is momentarily at rest at the upper turning point of its motion.

Example 38

(a) Show, by direct substitution, that the function $y(x,t) = A\sin(\omega t - kx)$ is a solution of the one-dimensional wave equation $\frac{\partial^2 y}{\partial t^2} = v^2\frac{\partial^2 y}{\partial x^2}$, and identify the wave speed v in terms of ω and k.

(b) A transverse wave on a stretched wire of linear mass density $\mu = 5.0\times10^{-3}\text{kg/m}$ is described by $y(x,t) = 0.010\sin(80\pi t - 8\pi x)\text{m}$, with x in metres and t in seconds. Find the tension in the wire and the maximum acceleration of any particle of the wire.

Solution

(a) Differentiate $y = A\sin(\omega t - kx)$ twice with respect to t, holding x fixed:

$$\frac{\partial y}{\partial t} = A\omega\cos(\omega t - kx), \qquad \frac{\partial^2 y}{\partial t^2} = -A\omega^2\sin(\omega t - kx)$$

Differentiate twice with respect to x, holding t fixed:

$$\frac{\partial y}{\partial x} = -Ak\cos(\omega t - kx), \qquad \frac{\partial^2 y}{\partial x^2} = -Ak^2\sin(\omega t - kx)$$

Substitute into the wave equation:

$$-A\omega^2\sin(\omega t - kx) = v^2(-Ak^2\sin(\omega t - kx))$$

The equality holds for all x and t provided $\omega^2 = v^2k^2$, that is

$$v = \frac{\omega}{k}$$

So $y = A\sin(\omega t - kx)$ is a solution of the wave equation with wave speed equal to the ratio of the angular frequency to the wave number.

(b) Comparing $y = 0.010\sin(80\pi t - 8\pi x)\text{m}$ with $y = A\sin(\omega t - kx)$:

$$A = 0.010\text{m}, \qquad \omega = 80\pi\text{rad/s}, \qquad k = 8\pi\text{rad/m}$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{80\pi\text{rad/s}}{8\pi\text{rad/m}} = 10\text{m/s}$$

Tension:

$$T = \mu v^2 = (5.0\times10^{-3}\text{kg/m})(10\text{m/s})^2 = 0.50\text{N}$$

Maximum particle acceleration:

$$a_{max} = A\omega^2 = (0.010\text{m})(80\pi\text{rad/s})^2 = (0.010)(6400\pi^2)\text{m/s}^2 = 64\pi^2\text{m/s}^2 \approx 632\text{m/s}^2$$

Example 39

(a) Account in physical terms for the fact that two particles separated by an integer number of wavelengths along a progressive wave are always in phase, regardless of the magnitude of the separation.

(b) A transverse wave of wavelength 1.20m and period 0.40s travels along a long stretched string. Find the phase difference between two points on the string separated by 0.30m at the same instant, and the phase difference between the displacements of the same point at two instants separated by 0.05s.

Solution

(a) The phase difference between two points separated by distance Δx on the same wave is $\Delta\varphi = k\,\Delta x = (2\pi/\lambda)\Delta x$. When $\Delta x = n\lambda$ for integer n, the phase difference is $\Delta\varphi = 2\pi n$, which is an integer number of full cycles. A sinusoidal function repeats exactly after every full cycle, so the two particles return to identical positions and identical velocities at every instant. The agreement is independent of the magnitude of Δx because only the fractional cycle counts; an integer number of wavelengths contributes no fractional cycle.

(b) **Spatial** phase difference:

$$\Delta\varphi = \frac{2\pi}{\lambda}\Delta x = \frac{2\pi}{1.20\text{m}}(0.30\text{m}) = \frac{\pi}{2}\text{rad}$$

Temporal phase difference:

$$\Delta\varphi = \frac{2\pi}{T}\Delta t = \frac{2\pi}{0.40\text{s}}(0.05\text{s}) = \frac{\pi}{4}\text{rad}$$

Example 40

(a) Account for the experimental observation that two coherent loudspeakers driven in phase produce regions in front of them that are permanently silent, regardless of the volume.

(b) Two coherent loudspeakers radiate notes of frequency 680Hz in still air, where the speed of sound is 340m/s. At a point P the path difference from the two loudspeakers is 0.25m. State, with reasons, whether P is a region of constructive or of destructive interference.

Solution

(a) Coherent sources radiate waves whose phase difference at any point depends only on the path difference from the two sources to that point. At a point where the path difference is an odd number of half-wavelengths, the two waves arrive in antiphase, the displacements cancel exactly at every instant, and the resultant amplitude is zero. The cancellation is an algebraic identity between two sinusoids of equal amplitude and opposite sign; it persists no matter how large the volume of either source, since louder sources cancel just as completely as faint ones.

(b) Wavelength of the radiated sound:

$$\lambda = \frac{v}{f} = \frac{340\text{m/s}}{680\text{Hz}} = 0.50\text{m}$$

Phase difference at P:

$$\Delta\varphi = \frac{2\pi}{\lambda}\Delta x = \frac{2\pi}{0.50\text{m}}(0.25\text{m}) = \pi\text{rad}$$

Since $\Delta\varphi = \pi\text{rad}$ (equivalently, the path difference is exactly half a wavelength), the two waves arrive at P exactly out of phase. Their displacements cancel at every instant, and P is a region of destructive interference, that is, a permanent silent region in front of the speakers.

Example 41

(a) Account for the fact that $y = A\sin(\omega t - kx)$ represents a wave travelling in the $+x$ direction, while $y = A\sin(\omega t + kx)$ represents one travelling in the $-x$ direction.

(b) A point source of acoustic power 12W radiates uniformly in still air. Find the intensity of the sound at 4.0m from the source, the distance from the source at which the intensity has fallen to 1.0mW/m^2, and the ratio of the displacement amplitudes at the two distances.

Solution

(a) A point of constant phase on the wave $y = A\sin(\omega t - kx)$ satisfies $\omega t - kx = \text{constant}$.

Differentiating with respect to time gives $\omega - k\frac{dx}{dt} = 0$, so:

$$\frac{dx}{dt} = \frac{\omega}{k} = +v$$

The point of constant phase therefore moves in the direction of increasing x: the wave pattern slides toward $+x$. Repeating the same calculation for $y = A\sin(\omega t + kx)$ gives $\frac{dx}{dt} = -\frac{\omega}{k} = -v$, so the wave pattern slides toward $-x$. The sign in front of kx is the wave's direction-of-travel indicator.

(b) Intensity at $r_1 = 4.0\text{m}$:

$$I_1 = \frac{P}{4\pi r_1^2} = \frac{12\text{W}}{4\pi(4.0\text{m})^2} = \frac{12\text{W}}{64\pi\text{m}^2} = 5.97\times10^{-2}\text{W/m}^2$$

Distance at which $I_2 = 1.0\times10^{-3}\text{W/m}^2$:

$$r_2^2 = \frac{P}{4\pi I_2} = \frac{12\text{W}}{4\pi(1.0\times10^{-3}\text{W/m}^2)} = 955\text{m}^2,$$

$$r_2 = \sqrt{955\text{m}^2} = 30.9\text{m}$$

Since $I \propto A^2$, $A \propto 1/r$, hence

$$\frac{A_1}{A_2} = \frac{r_2}{r_1} = \frac{30.9\text{m}}{4.0\text{m}} = 7.7$$

The amplitude at 4.0m is about 7.7 times the amplitude at 30.9m.

Example 42

(a) Account for the observation that, when two stereo loudspeakers play the same note in phase, a listener positioned exactly between them hears a louder note than a listener standing close to either speaker alone.

(b) Two waves of equal frequency and amplitudes 6.0cm and 4.0cm arrive at the same point with a phase difference of $\frac{2\pi}{3}\text{rad}$. Find the amplitude of the resultant.

Solution

(a) A listener exactly between the two coherent speakers receives the two waves with zero path difference and therefore zero phase difference. The two displacements add at every instant, giving a resultant of amplitude A_1+A_2, which is twice the amplitude received from either speaker alone (when the two amplitudes are equal). Since intensity goes as the square of

amplitude, the listener at the midpoint perceives roughly four times the intensity of one speaker alone, which is appreciably louder. A listener close to one speaker is essentially in the field of that speaker only and receives a single amplitude A_1.

(b) For two waves of amplitudes $A_1 = 6.0\text{cm}$ and $A_2 = 4.0\text{cm}$ with phase difference $\Delta\varphi = \frac{2\pi}{3}\text{rad}$, the resultant amplitude is given by:

$$A_R^2 = A_1^2 + A_2^2 + 2A_1A_2\cos(\Delta\varphi)$$

Substituting:

$$A_R^2 = (6.0\text{cm})^2 + (4.0\text{cm})^2 + 2(6.0\text{cm})(4.0\text{cm})\cos\frac{2\pi}{3}$$

$$A_R^2 = 36\text{cm}^2 + 16\text{cm}^2 + 2(24\text{cm}^2)\left(-\frac{1}{2}\right) = 52\text{cm}^2 - 24\text{cm}^2 = 28\text{cm}^2$$

$$A_R = \sqrt{28\text{cm}^2} = 2\sqrt{7}\text{cm} \approx 5.3\text{cm}$$

Example 43

(a) Account for the fact that two musical notes of slightly different frequency, when sounded together, are heard as a single note whose loudness rises and falls in a slow, regular rhythm.

(b) A piano tuner sounds a tuning fork of frequency 256Hz together with one of the strings of a piano and counts 12 beats in 4s. He then loads the string very lightly with a small piece of wax and the beat rate falls to 2Hz. Find the original frequency of the piano string.

Solution

(a) When two close-frequency notes sound together their displacements at the listener's ear superpose. The two waves drift slowly into and out of phase: at moments when they are in phase the resultant amplitude is $A_1 + A_2$ (loud) and at moments when they are in antiphase the resultant amplitude is $|A_1 - A_2|$ (faint or silent). The cycle of loudness repeats at a frequency equal to the difference of the two source frequencies, $f_b = |f_1 - f_2|$. These regular pulsations of loudness are the beats.

(b) Initial beat frequency:

$$f_b = \frac{12}{4\text{s}} = 3\text{Hz}$$

The piano string is therefore at $f_{\text{string}} = 256\text{Hz} \pm 3\text{Hz}$, i.e., 253Hz or 259Hz.

Loading the string with wax increases its linear mass density. From the third law of vibration of a stretched string, $f \propto 1/\sqrt{\mu}$, so loading lowers the string's frequency. After waxing the beat rate has fallen from 3Hz to 2Hz, so the string's frequency must have moved closer to the fork.

If the original frequency was 259Hz, lowering it slightly produces a string at 258Hz and a beat of 2Hz, which is consistent with observation.

If the original frequency was 253Hz, lowering it would produce a string at 252Hz or below and a beat of 4Hz or more, which is inconsistent with observation.

The original frequency of the piano string is $f_{\text{string}} = 259\text{Hz}$.

Example 44

(a) Account for the experimental observation that loading a slightly mistuned guitar string with a small piece of plasticine can either remove beats with a reference fork or make them faster, depending on which side of the fork frequency the string was on.

(b) A sonometer wire of length 25.0cm is stretched between two bridges under fixed tension. When the wire is sounded together with a tuning fork of frequency 384Hz, five beats per second are heard. The wire is shortened, with no change in tension, until the beats vanish. Find the new length of the wire.

Solution

(a) The beat frequency between the fork and the string is $f_b = |f_{\text{fork}} - f_{\text{string}}|$. Loading the string with plasticine increases its mass per unit length, so by the third law of vibration its frequency falls. If the string was originally above the fork, lowering it reduces $|f_{\text{fork}} - f_{\text{string}}|$, so the beat frequency falls and may pass through zero (silence) before rising again. If the string was originally below the fork, lowering it increases $|f_{\text{fork}} - f_{\text{string}}|$, so the beat frequency rises immediately. The direction of the change in beat rate therefore reveals on which side of the fork the string was tuned.

(b) Initial beat frequency $f_b = 5\text{Hz}$, so the original string frequency is

$$f_{\text{old}} = 384\text{Hz} \pm 5\text{Hz} = 379\text{Hz or } 389\text{Hz}.$$

Shortening the wire raises its frequency (first law of vibration, $f \propto 1/L$). The shortening removes the beats, i.e., it raises the string to 384Hz. The original frequency must therefore have been below 384Hz, namely $f_{\text{old}} = 379\text{Hz}$.

Applying $f \propto 1/L$ at constant tension and mass per unit length:

$$\frac{L_{new}}{L_{old}} = \frac{f_{old}}{f_{new}} = \frac{379\text{Hz}}{384\text{Hz}},$$

$$L_{new} = (25.0\text{cm})\left(\frac{379\text{Hz}}{384\text{Hz}}\right) = 24.67\text{cm}$$

The wire must be shortened to 24.67cm for the beats with the fork to vanish.

Example 45

(a) Account in physical terms for the fact that a stretched string fixed at both ends vibrates only at a discrete set of frequencies, and not at any frequency intermediate between two consecutive members of the set.

(b) A string of length 1.20m fixed at both ends has linear mass density $\mu = 4.0 \times 10^{-3}\text{kg/m}$. Find the fundamental frequency of the string when the tension is 250N, and find the new tension required for the string to resonate at 400Hz in its fundamental mode.

Solution

(a) The transverse displacement at each fixed end is permanently zero, so any standing wave on the string must have a node at every support. The simplest such wave fits half a wavelength between the supports; the next fits one full wavelength; the next fits one-and-a-half; and so on. Only those wavelengths satisfying $L = n\lambda/2$ for positive integer n produce a stationary pattern that is consistent with both boundary conditions. The corresponding frequencies $f_n = nv/(2L)$ form a discrete set; intermediate frequencies cannot establish a self-consistent stationary pattern and decay quickly.

(b) Wave speed at the original tension:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{250\text{N}}{4.0 \times 10^{-3}\text{kg/m}}} = 250\text{m/s}$$

Fundamental frequency:

$$f_1 = \frac{v}{2L} = \frac{250\text{m/s}}{2(1.20\text{m})} = 104.2\text{Hz}$$

From the second law of vibration, $f_1 \propto \sqrt{T}$ at fixed L and μ, so

$$\frac{T_{new}}{T_{old}} = \left(\frac{f_{new}}{f_{old}}\right)^2 = \left(\frac{400\text{Hz}}{104.2\text{Hz}}\right)^2 = 14.74$$

$$T_{new} = 14.74 \times 250\text{N} = 3.69 \times 10^3\text{N}.$$

Example 46

(a) Account for the observation that bowing a sonometer wire at exactly its midpoint excites a much purer note than bowing it close to one of the bridges.

(b) A sonometer wire of length 1.00m fixed at both ends is excited in its third harmonic mode. Three small paper riders rest on the wire at distances 0.250m, 0.333m and 0.500m from one of the bridges. State, with reasons, which of the riders are thrown off the wire and which remain at rest.

Solution

(a) A vibrating fixed-fixed wire can support only those harmonics that have a node at every point at which the wire is forced. The midpoint $L/2$ is a node for every even harmonic ($n = 2,4,6, ...$) and an antinode for every odd one. Bowing the midpoint forces it to oscillate, which is incompatible with a node there, so every even harmonic is suppressed and only the odd ones survive. The lowest odd harmonic, namely the fundamental, dominates and the sound is comparatively pure. Bowing close to one bridge, where the wire is almost a node of the fundamental but almost an antinode of every higher harmonic, excites all the harmonics, and the resulting note carries a rich complement of overtones.

(b) For the third harmonic on a string of length $L = 1.00\text{m}$ fixed at both ends:

$$\lambda = \frac{2L}{3} = \frac{2(1.00\text{m})}{3} = 0.667\text{m}.$$

Nodes occur at $x = 0, \frac{L}{3}, \frac{2L}{3}$, $L = 0,\ 0.333\text{m},\ 0.667\text{m},\ 1.000\text{m}$.

Antinodes occur at $x = \frac{L}{6}, \frac{L}{2}, \frac{5L}{6} = 0.167\text{m},\ 0.500\text{m},\ 0.833\text{m}$.

Rider at $x = 0.250\text{m}$: lies between a node at 0 and an antinode at 0.167m, so the wire there has a non-zero displacement. The rider is thrown off.

Rider at $x = 0.333\text{m}$: lies exactly on a node, so the wire at this point does not move at all. The rider remains at rest.

Rider at $x = 0.500\text{m}$: lies exactly on an antinode, where the wire has its largest displacement. The rider is thrown off vigorously.

Example 47

(a) Account in physical terms for the experimental observation that two steel piano wires of the same length, material and tension produce different fundamental notes when their diameters differ.

(b) A steel piano wire of diameter 0.80mm and tension 600N produces a fundamental note of frequency 220Hz. The wire is replaced by a second steel wire of the same length under the same tension but of diameter 0.96mm. Find the fundamental frequency of the new wire.

Solution

(a) The fundamental frequency of a stretched wire is $f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$, where the linear mass density of a uniform wire of circular cross-section is $\mu = \rho\pi d^2/4$. For wires of the same material, ρ is fixed and a thicker wire has a larger μ. At equal length and tension, a larger μ gives a smaller wave speed $v = \sqrt{T/\mu}$ and therefore a smaller fundamental frequency. Two wires that are otherwise identical sound at different pitches purely because their diameters set different masses per unit length.

(b) For wires of the same material, $\mu \propto d^2$. At constant tension and length, $f_1 \propto 1/\sqrt{\mu} \propto 1/d$. Hence

$$\frac{f_{new}}{f_{old}} = \frac{d_{old}}{d_{new}} = \frac{0.80\text{mm}}{0.96\text{mm}} = \frac{5}{6}$$

$$f_{new} = \frac{5}{6}(220\text{Hz}) = 183.3\text{Hz}$$

Example 48

(a) Distinguish between a progressive wave and a stationary wave with respect to the transport of energy along the medium and the variation of amplitude with position along the medium.

(b) The displacement of a stationary wave on a stretched string is given by $y(x,t) = 0.040\sin(5\pi x)\cos(60\pi t)\text{m}$, with x in metres and t in seconds. Find the amplitude of each constituent travelling wave, the wavelength, the speed of those travelling waves, the position of the antinode nearest to $x = 0$ and the displacement of that antinode at the instant $t = 0.025\text{s}$.

Solution

(a) A progressive wave transports energy continuously through the medium in the direction in which it travels; a stationary wave does not transport any net energy along the medium, energy being exchanged locally between kinetic and potential forms in each cell between consecutive nodes.

Every particle of a progressive wave oscillates with the same amplitude as every other particle. In a stationary wave the amplitude varies sinusoidally with position: it is zero at the nodes (spaced $\lambda/2$ apart) and is a maximum equal to twice the amplitude of either constituent travelling wave at the antinodes (located midway between consecutive nodes).

(b) Compare $y = 0.040\sin(5\pi x)\cos(60\pi t)\text{m}$ with $y = 2A\sin(kx)\cos(\omega t)$:

$$2A = 0.040\text{m}, \qquad k = 5\pi\text{rad/m}, \qquad \omega = 60\pi\text{rad/s}.$$

Amplitude of each constituent travelling wave:

$$A = \frac{0.040\text{m}}{2} = 0.020\text{m}$$

Wavelength:

$$\lambda = \frac{2\pi}{k} = \frac{2\pi}{5\pi\text{rad/m}} = 0.40\text{m}$$

Wave speed:

$$v = \frac{\omega}{k} = \frac{60\pi\text{rad/s}}{5\pi\text{rad/m}} = 12\text{m/s}$$

Antinodes are where $\sin(kx) = \pm 1$, i.e., $kx = (2n+1)\pi/2$ for integer $n \geq 0$. The nearest to $x = 0$ corresponds to $n = 0$:

$$x = \frac{\pi/2}{k} = \frac{\pi/2}{5\pi\text{rad/m}} = 0.10\text{m}$$

Displacement of that antinode at $t = 0.025\text{s}$:

$$y = (0.040\text{m})\sin(5\pi \times 0.10)\cos(60\pi \times 0.025) = (0.040\text{m})\sin\frac{\pi}{2}\cos\frac{3\pi}{2} = (0.040\text{m})(1)(0) = 0.$$

The antinode is momentarily passing through its equilibrium position, where it is moving with maximum speed.

Example 49

(a) A steel wire on a stringed instrument tuned in a cool morning sounds noticeably flat by midday, even though no physical alteration is made to the instrument. Account for this observation.

(b) A steel wire of length 0.65m, cross-sectional area $8.0 \times 10^{-7}\text{m}^2$, density 7800kg/m^3, Young's modulus $2.0 \times 10^{11}\text{N/m}^2$ and linear thermal-expansion coefficient $\alpha = 1.2 \times 10^{-5}\text{K}^{-1}$ is rigidly clamped between two fixed

supports and tuned at $20°C$ to a fundamental frequency of $200Hz$. Find the fundamental frequency of the wire when its temperature has risen to $35°C$.

Solution

(a) When the temperature rises, the wire tries to expand longitudinally because of thermal expansion. The supports are rigid and prevent that expansion, so the tension that the supports must apply to keep the wire of unchanged length is reduced by an amount $\Delta T = EA\alpha\,\Delta\theta$. The fundamental frequency of a stretched wire varies as $f_1 \propto \sqrt{T}$, so a smaller tension gives a smaller fundamental frequency: the wire goes flat.

(b) Linear mass density:

$$\mu = \rho A = (7800 kg/m^3)(8.0 \times 10^{-7} m^2) = 6.24 \times 10^{-3} kg/m$$

Initial tension from $f_1 = \frac{1}{2L}\sqrt{\frac{T_1}{\mu}}$:

$$T_1 = 4\mu L^2 f_1^2 = 4(6.24 \times 10^{-3} kg/m)(0.65m)^2(200Hz)^2 = 421.8N$$

Drop in tension caused by the temperature rise of $\Delta\theta = 35°C - 20°C = 15K$:

$$\Delta T = EA\alpha\,\Delta\theta = (2.0 \times 10^{11} N/m^2)(8.0 \times 10^{-7} m^2)(1.2 \times 10^{-5} K^{-1})(15K) = 28.8N$$

New tension:

$$T_2 = T_1 - \Delta T = 421.8N - 28.8N = 393.0N$$

New fundamental frequency:

$$f_2 = f_1\sqrt{\frac{T_2}{T_1}} = (200Hz)\sqrt{\frac{393.0N}{421.8N}} \approx 193Hz$$

The wire has dropped from $200Hz$ to about $193Hz$, a flattening of roughly $7Hz$, in agreement with the qualitative observation in part (a).

Example 50

(a) Account for the fact that a transverse pulse generated at the lower (free) end of a long heavy hanging rope speeds up as it travels upward toward the fixed support.

(b) A uniform rope of length $5.0m$ and mass $2.0kg$ hangs vertically from a fixed support, with its lower end free. Find the time taken for a small transverse pulse to travel from the free lower end to the support. Take $g = 9.8 m/s^2$.

Solution

(a) Each element of the rope must support the weight of the rope hanging beneath it. At a height y above the free end, the tension equals the weight of the segment of length y:

$$T(y) = \mu g y,$$

where μ is the linear mass density of the rope. The transverse wave speed at that height is

$$v(y) = \sqrt{\frac{T(y)}{\mu}} = \sqrt{gy},$$

which depends only on g and on the height y above the free end. As the pulse climbs, y grows and so does $v(y)$: the pulse speeds up because the tension supporting it grows.

(b) The pulse travels with the position-dependent speed $v(y) = \sqrt{gy}$. The time to travel from $y = 0$ to $y = L$ is

$$t = \int_0^L \frac{dy}{v(y)} = \frac{1}{\sqrt{g}}\int_0^L y^{-1/2}\,dy = \frac{1}{\sqrt{g}}\left[2\sqrt{y}\right]_0^L = 2\sqrt{\frac{L}{g}}$$

Substituting $L = 5.0m$ and $g = 9.8 m/s^2$:

$$t = 2\sqrt{\frac{5.0m}{9.8 m/s^2}} = 1.43s$$

The pulse takes about $1.43s$ to climb the rope from the free end to the support.

Example 51

(a) Account for the fact that every particle of the medium through which a sinusoidal progressive wave passes executes simple harmonic motion about its undisturbed position.

(b) A wave on a stretched string is given by $y(x,t) = 0.020\sin(40\pi t - \pi x)m$, with x in metres and t in seconds. For the particle at $x = 0.50m$, find the period of its motion, the maximum speed, the maximum acceleration, and the total mechanical energy per unit mass.

Solution

(a) For a sinusoidal travelling wave $y = A\sin(\omega t - kx)$, the displacement of the particle at any fixed position x_0 is $y(t) = A\sin(\omega t - kx_0)$. The acceleration of that particle is

$$\frac{\partial^2 y}{\partial t^2} = -A\omega^2 \sin(\omega t - kx_0) = -\omega^2 y$$

This is the defining equation of simple harmonic motion: acceleration directed toward, and proportional to, the displacement from equilibrium. Every particle of the medium is therefore a simple harmonic oscillator with the same angular frequency ω and the same amplitude A as every other particle, differing only in the constant phase shift $-kx_0$.

(b) From $y = 0.020\sin(40\pi t - \pi x)\text{m}$ we read off amplitude $A = 0.020\text{m}$ and angular frequency $\omega = 40\pi\text{rad/s}$, with a constant phase shift at $x = 0.50\text{m}$ that does not change any of the quantities asked for.
Period:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{40\pi\text{rad/s}} = \frac{1}{20}\text{s} = 0.050\text{s}$$

Maximum speed:

$$v_{p,max} = A\omega = (0.020\text{m})(40\pi\text{rad/s}) = 0.8\pi\text{m/s} \approx 2.51\text{m/s}$$

Maximum acceleration:

$$a_{max} = A\omega^2 = (0.020\text{m})(40\pi\text{rad/s})^2 = (0.020)(1600\pi^2)\text{m/s}^2 = 32\pi^2\text{m/s}^2 \approx 316\text{m/s}^2$$

Total mechanical energy per unit mass of the particle:

$$\frac{E}{m} = \frac{1}{2}\omega^2 A^2 = \frac{1}{2}(40\pi\text{rad/s})^2(0.020\text{m})^2 = \frac{1}{2}(1600\pi^2)(4.0\times10^{-4})\text{J/kg} = 0.32\pi^2\text{J/kg} \approx 3.16\text{J/kg}$$

Example 52

(a) Account in physical terms for the experimental observation that a soprano holding a single sustained note close to a wine glass can shatter the glass, while singing a different note at the same volume close to the same glass produces no effect.
(b) A wine glass has a natural mechanical frequency of 720Hz when struck. The minimum sound intensity at the glass required to shatter it is 12W/m^2. A soprano's voice radiates 0.020W of acoustic power uniformly in all directions. Find the maximum distance from her mouth at which the glass can be shattered.

Solution

(a) A wine glass is a lightly damped mechanical oscillator with a sharp natural frequency, set by its mass and stiffness. When the singer holds a note exactly equal to that natural frequency, the glass is driven at resonance: each cycle of the sound wave delivers a small additional store of energy in step with the glass's own oscillation, and the amplitude of the glass's flexure builds up cycle by cycle until the strain in the glass exceeds its breaking strength. When the singer sings a different note the driving frequency is detuned from the natural frequency; the energy delivered per cycle is on average small (some cycles in phase, some out), the response amplitude saturates at a small value, and the glass merely vibrates faintly.

(b) For a point source the intensity at a distance r is

$$I = \frac{P}{4\pi r^2}$$

Setting I equal to the shattering threshold:

$$r^2 = \frac{P}{4\pi I} = \frac{0.020\text{W}}{4\pi(12\text{W/m}^2)} = 1.33\times10^{-4}\text{m}^2$$

$$r = \sqrt{1.33\times10^{-4}\text{m}^2} = 0.0115\text{m} = 1.15\text{cm}$$

The glass can be shattered only if it is held within about 1.15cm of the singer's mouth, which is why operatic glass-shattering demonstrations are always performed with the glass held very close to the speaker.

With twenty miscellaneous worked examples now complete, the picture should feel nicely connected. It is time to stretch our understanding a little further. Welcome to the Digging Deeper Exercise!

DIGGING DEEPER EXERCISE 7

EXERCISE 7A: BINDER QUESTIONS

Question 1

Account in physical terms for the fact that the speed at which a progressive wave travels through a given medium does not depend on the amplitude of the wave.

Question 2

Kipanga writes the maximum speed of a particle of a transverse wave as $v_{p,max} = A\omega$, while Kipute writes it as $v_{p,max} = 2\pi fA$. Each student insists that the two formulas describe different physical quantities. Account for the fact that the two formulas are equivalent.

Question 3

Account in physical terms for the fact that the frequency of a progressive wave is preserved when the wave crosses a boundary between two different media.

Question 4

Show that the displacement amplitude of a wave radiating uniformly outward from a point source in open space falls as $1/r$ at a distance r from the source.

Question 5

Two coherent sources radiating into a hall produce regions of permanent silence by destructive interference. Account in physical terms for the fate of the energy delivered by the two sources at and around such a silent point.

Question 6

A student claims that the formula $A_R = 2A\cos(\Delta\varphi/2)$ for the resultant amplitude of two superposing waves of the same frequency holds only when the two waves have equal amplitudes. Account in physical terms for this restriction.

Question 7

Classify each of the wave equations $y_1 = A\sin(\omega t - kx)$, $y_2 = 2A\sin(kx)\cos(\omega t)$ and $y_3 = A\sin(\omega t + kx)$ as describing a travelling or a stationary wave, giving reasons in each case.

Question 8

Two students argue about the energy of a stationary wave on a stretched string fixed at both ends. Kipanga claims that, since a stationary wave does not transport energy along the string, the string carries no kinetic energy. Kipute claims that every particle of the string is in motion (except at the nodes), so kinetic energy must be present. Resolve the disagreement.

Question 9

Account in physical terms for the experimental observation that the fundamental frequency of a stretched string fixed at both ends decreases as the linear mass density of the string is increased, with the length of the string and the tension in it both held constant.

Question 10

A student observing a stretched string vibrating in its third harmonic claims that, since the frequency is now three times that of the fundamental, every particle of the string must be moving three times as fast as it would be moving if the string were vibrating in the fundamental at the same antinode amplitude. Indicate, with reasons, whether this claim is correct.

Question 11

Explain in physical terms why a thicker uniform metal rod, made of the same material as a thinner one, transmits a longitudinal wave at the same speed as the thinner rod.

EXERCISE 7B: REAL QUESTIONS

Question 12

A leaf placed on the calm surface of Lake Victoria, when a fisherman drops a stone a metre away from it, is observed to bob up and down repeatedly without travelling along with the ripples that pass beneath it. Account in physical terms for what is moving and what is not.

Question 13

A radiologist at the Bugando Hospital in Mwanza directs a 5MHz ultrasound pulse from the soft tissue of a patient's abdomen into the underlying muscle, where the wave travels at a higher speed than in fat. Account in physical terms for the changes in the wave's parameters at the boundary.

Question 14

A villager in Shinyanga, awaiting the arrival of his herd in the evening, presses his ear to the dry ground and hears the rumble of the cattle hooves long before the sound reaches him through the open air. Account in physical terms for why pressing the ear to the ground gives an earlier warning.

Question 15

Two coherent loudspeakers, mounted at opposite ends of the stage of the Karimjee Hall in Dar es Salaam and driven from the same amplifier at 680Hz, produce a steady note. A patron walking slowly along a line parallel to the stage notices that the loudness rises and falls at regular intervals. Account in physical terms for the variation.

Question 16

A guitarist tuning his instrument at a wedding rehearsal in Unguja sounds his fifth string together with a 440Hz reference fork and counts six beats per second. He tightens the string slightly and the beat rate falls to two beats per second. Identify, with reasons, the original frequency of the fifth string.

Question 17

A guitarist at a school concert in Mtwara presses a finger firmly against the twelfth fret of his guitar string and finds that the resulting note sounds exactly an octave above the open-string note. Deduce the position of the twelfth fret along the vibrating length of the string.

Question 18

A student in Chakechake, watching the cables of an electricity pylon swing sideways during a strong wind, notices that one cable in particular vibrates with very large amplitude at one specific wind speed but is comparatively still at all other wind speeds. Account in physical terms for the speed-specific vibration.

Question 19

A washing-line stretched tightly between two posts in a courtyard at Songea, when struck with a stick, sends a sharp transverse pulse to the other post in a fraction of a second. The same line, when its tension is reduced to a slack droop, sends a slow visible wave when struck. Account in physical terms for the difference in propagation speed.

Question 20

Two tuning forks, both nominally of frequency 440Hz, are sounded together in a music classroom in Tanga. Instead of a steady single note, the class hears a slow rhythmic rise and fall in loudness, repeating roughly twice each second. Account in physical terms for this rhythm.

Question 21

A pedestrian crossing the Bridge during a strong wind notices that the long stay-cables sing with a clearly audible note when the wind reaches one particular speed but are silent above and below it. Account in physical terms for why the singing is heard only at one particular wind speed.

EXERCISE 7C: HOT QUESTIONS

Question 22

A transverse wave on a stretched string is described by $y(x,t) = 0.060\sin(50\pi t - 2\pi x)\text{m}$, with x in metres and t in seconds. Determine the wavelength, the period and the wave speed, and calculate the maximum speed and the maximum acceleration of any particle of the string.

Question 23

A point source radiates sound at an acoustic power of 4.0W uniformly in still air. Calculate the intensity of the sound at distances of 5.0m and 10.0m from the source, and the ratio of the displacement amplitudes at the two distances.

Question 24

Two waves of equal amplitude 4.0cm travel in the same direction along a stretched rope. Their displacements at a fixed point are $y_1 = 4.0\sin(\omega t + 5\pi/6)\text{cm}$ and $y_2 = 4.0\sin(\omega t + \pi/6)\text{cm}$. Determine the amplitude of the resultant and the phase angle of the resultant relative to y_1.

Question 25

Two coherent loudspeakers are placed 1.50m apart and emit notes of frequency 680Hz in still air, where the speed of sound is 340m/s. A point P is observed at a distance of 4.00m from the first loudspeaker and 5.50m from the second. Compute the wavelength of the sound and the path difference at P, and state, with reasons, whether P lies in a region of constructive or of destructive interference.

Question 26

A tuning fork of frequency 512Hz is sounded together with a stretched sonometer wire under a tension of 90N. Eight beats per second are heard. The tension is then increased until the beats vanish. Determine, with reasons, the original frequency of the wire, and calculate the new tension at which the beats vanish.

Question 27

A stretched wire of length 0.80m, fixed at both ends and of linear mass density $\mu = 5.0 \times 10^{-3}\text{kg/m}$, is observed to vibrate with two antinodes between its supports at a frequency of 240Hz. Calculate the speed of a transverse wave along the wire, and the tension in the wire.

Question 28

A stretched wire of length 1.20m, fixed at both ends and held under a tension of 200N, has linear mass density $5.0 \times 10^{-3}\text{kg/m}$. Determine the wave speed and the fundamental frequency, and find the highest harmonic of the wire that can be heard by a listener whose upper limit of audible frequency is 20kHz.

Question 29

A copper wire has tensile strength $\sigma_{\text{max}} = 2.2 \times 10^{8}\text{N/m}^2$ and density $\rho = 8900\text{kg/m}^3$. Show that the maximum speed at which a transverse wave can travel along the wire depends only on σ_{max} and ρ, and use the result to compute the maximum wave speed and the maximum tension that may be applied to a wire of cross-sectional area $1.5 \times 10^{-6}\text{m}^2$.

Question 30

An aluminium rod of length 50m has Young's modulus $Y = 7.0 \times 10^{10}\text{N/m}^2$ and density $\rho = 2700\text{kg/m}^3$. Calculate the speed of a longitudinal wave along the rod and the time taken by a pulse to travel from one end of the rod to the other.

Question 31

A stationary wave on a stretched string is given by $y(x,t) = 0.030\sin(4\pi x)\cos(80\pi t)\text{m}$, with x in metres and t in seconds. Determine the amplitude of each constituent travelling wave, the wavelength, the wave speed, and the position of the antinode nearest to $x = 0$. Calculate also the displacement of that antinode at the instant $t = 0.020\text{s}$.

Chapter 8

PHYSICAL OPTICS

INTRODUCTION

The morning assembly at Miono Secondary had run a little long because the head teacher had a great many announcements to make. By the time it was over the sun had moved high enough to send a sharp white shaft through the louvre windows and on to the back row of desks. Kipanga had brought a CD to school for reasons he could no longer remember, and the disc was lying face-up on the wooden top in front of him. As the light struck the reflective surface, the CD did something the sun, by itself, had never been able to do. It threw a fan of colour across the desk: red on the outside, then orange, yellow, green, blue, and a violet so faint it could only be seen out of the corner of one eye. The colours were not painted on the disc. They had not been there a moment before. They appeared with the light, shifted with the disc, and disappeared the moment Kipanga moved his head.

He stared at the colours for long enough that Kipute, sitting next to him, leaned across to see what was so interesting. Then Mr. Akilikubwa, who had been standing at the doorway of the form-five classroom waiting for the assembly to disperse, spotted the boys hunched over the desk and walked across without comment.

Kipanga: *"It looks like a prism."*

Mr. Akilikubwa: *"It is not a prism, Kipanga."*

Kipanga: *"But it is splitting white light into colours."*

Mr. Akilikubwa: *"It is doing exactly that. It is not, however, doing it the way a prism does it. A prism splits white light by refraction; the speed of light in glass depends on the colour, so different colours bend by different amounts. There is no glass here, and there is no refraction. What is doing the work is the surface of the disc itself, which has been ruled with about six hundred twenty-five microscopic tracks per millimetre during manufacture. You are holding a diffraction grating, and you have been carrying it in your school bag all term without knowing it. By the time this chapter is over, you will be able to compute the angle at which red comes out, and the angle at which violet comes out, and the spacing between successive tracks on this disc, all by ear and entirely from your head."*

He paused, picked up the disc, and tilted it under the light.

Mr. Akilikubwa: *"This is the first of three phenomena we are going to spend a chapter on. The second is the reason a soap bubble shows rainbow colours just before it bursts, and the reason a thin film of petrol left on a Mwanza puddle does the same a few minutes after the rain stops. The third is the reason your aunt's polaroid sunglasses kill the glare from the surface of Lake Victoria, and the reason an ordinary pair only darkens it. All three are wave phenomena, and all three need an explanation that the ray model of light, useful as it has been to us so far, simply cannot supply. Welcome to the part of optics where light remembers it is a wave."*

Welcome to Chapter 8. The previous chapter taught you what a wave is, what a wave equation looks like, and what happens when two waves meet at a point. This chapter takes those three pieces of knowledge and applies them, one after the other, to light. The result is a small library of phenomena that the ray optics of your earlier school years could not even pose, let alone solve. We are about to give a wave-mechanical account of why a CD shows rainbows; why a soap bubble displays a band of colour just before it bursts; why two slits in a screen, illuminated by a single colour of light, produce a striped pattern of bright and dark fringes on a wall behind them; why a single slit, illuminated by the same light, produces a paler, broader, but still striped pattern of its own; why a curved lens placed on a flat glass plate shows a target of dark and bright concentric rings under monochromatic light; and why some of the light reflected off a window pane is polarised while the light coming directly from the sun above is not.

These phenomena go by three collective names. Two waves of light, taking different paths from the same source and meeting at the same point on a screen, produce **interference**. A single wave of light, passing an obstacle whose size is comparable with its wavelength, produces **diffraction**. And the underlying transverse character of light, which the previous chapter mentioned and the next chapter will exploit, has the consequence that the orientation of the electric field along the direction of propagation can be selected; a wave whose electric field lives in a single plane is called **plane-polarised**, and the process of making such a wave is called **polarisation**. The chapter is divided into three blocks, one for each phenomenon, plus a closing block on what physical optics is good for in the world.

Three pieces of Chapter 7 we will need over and over. The first is the principle of superposition: when two waves arrive at the same point of the medium, the resultant displacement at that point is the algebraic sum of the displacements that each wave would produce on its own. The second is the relation between path difference and phase difference, namely that a path difference Δx, for waves of wavelength λ, corresponds to a phase difference $\Delta\varphi$ given by

$$\Delta\varphi = \frac{2\pi}{\lambda}\Delta x$$

This is the equation we will use to predict where two waves will reinforce and where they will cancel. The third is the inverse-square law for the intensity of a spherical wave, which we will need when we ask why light from independent lamps cannot give sustained interference, and why a laser pointer can.

Two ideas are entirely new in this chapter and have to be installed before the rest will work. The first is **coherence**, *a property of two sources that determines whether their waves can produce a stable interference pattern at all.* Two sources are coherent if they maintain a fixed phase relationship. Two sources are incoherent if their relative phase wanders around the cycle on a timescale much shorter than your eye can follow. Most natural light sources are incoherent, which is the reason interference fringes are not the first thing you see when you walk into a sunlit room. The trick of every interference experiment in this chapter is to take a single coherent source and split it into two paths, so that whatever phase the source happens to be in at any instant, both paths share it. The two paths then recombine at a screen, and the only thing that determines where the fringes fall is the geometry of the apparatus.

The second new idea is the **wavefront**, *the locus of points in identical phase,* together with the **wavelet**, *the secondary spherical wave that every point on a wavefront sends out.* **Huygens' principle**, which we are about to develop in the very next section, says that *the next wavefront is the envelope of all those wavelets.* With wavefronts and wavelets in hand we can derive the law of reflection and the law of refraction without using any rays at all, which is a satisfying proof that the wave model contains the ray model, rather than the other way around. Wavelets are also the foundation of every diffraction calculation later in the chapter.

A short warning before we start. This chapter contains four equations that look almost identical and mean different things, namely:

$$\mathbf{y_n = \frac{n\lambda D}{d}}$$

Young's slits, position of the n-th bright fringe.

$$\mathbf{2t = n\lambda}$$

Air wedge, condition for the n-th dark fringe.

$$\mathbf{D_n^2 = 4n\lambda a}$$

Newton's rings, diameter of the n-th dark ring.

$$\mathbf{d\sin\theta = n\lambda}$$

Diffraction grating, n-th principal maximum.

The first describes positions of bright spots on a screen behind two slits. The second describes the thickness of an air film at which destructive interference occurs in reflected light. The third describes the diameters of dark rings between a curved lens and a flat plate. The fourth describes the angles at which a diffraction grating sends its bright spectra. They share an integer n and a wavelength λ, and they all follow from the same path-difference logic, which is why they look alike on the page. Telling them apart is two thirds of the chapter's arithmetic, and the rest is doing the algebra after you have decided which one applies. We will spend a little time in each section showing why each equation is the right one for its own geometry, and where the others would lead a student astray if borrowed by accident.

One last orientation. **Optics** is the part of physics where the wave model of light buys back its admission ticket. In your O-level physics studies, when you traced rays through lenses and mirrors and built images on paper with a ruler, you got remarkably far without ever asking what light is. By the end of this chapter you will know why that worked, and where it stops working. It works whenever the apertures and the obstacles are large compared with λ, the wavelength of light, which for visible light lies between roughly four hundred nanometres and seven hundred nanometres. It stops working the moment the apertures and the obstacles shrink to that scale, which is the moment interference and diffraction take over and the simple straight-line picture has to give way to something richer. The richer picture is what the rest of this chapter delivers. It is also why, in many of the questions you will meet in your examinations, the items that draw most heavily on this chapter are also the items that pay the highest mark per minute of work invested. Read this chapter once. Then read it again with a pencil. Then put a CD on a desk under the school light and check, with your own eyes, that the physics is true. Let us begin.

WAVEFRONTS AND HUYGENS' PRINCIPLE

Picture for a moment the surface of Lake Victoria after Mr. Akilikubwa's stone has dropped in. The ripples spread out as widening circles, neatly evenly spaced, all centred on the place where the stone hit the water. From above, what you see is not a single moving thing but a family of expanding circles. Now look at one of those circles, just one, and notice that every point on it is doing the same thing at the same instant. A crest is everywhere on the same circle; a trough is everywhere on the next-but-one circle, half a wavelength behind it; the place where the water surface is rising at exactly the still-water level is a third circle a quarter of a wavelength further behind. Each of these surfaces is a snapshot of one phase, captured around the lake at one instant. Each is what we are about to call a wavefront.

A **wavefront** *is the locus of all points of a wave that are in identical phase at the same instant.* The crest is a wavefront; the trough is a wavefront; the rising-zero-crossing is a wavefront. There are infinitely many wavefronts inside any wave, packed one against the next like layers of an onion. We usually draw only those that are exactly one wavelength apart, which gives

the picture you have in your head from O-level: a series of equally spaced lines or rings, each one wavelength apart from the next.

Two pictures of wavefronts will dominate the rest of this chapter. The first is the **plane wavefront**: a flat surface, like a clean sheet of A4 paper, advancing forward without changing shape. Light arriving from the sun, by the time it has crossed 150 million kilometres to your eyes, has wavefronts that are essentially planes; the source is so far away that you cannot detect the curvature of the wavefronts on a benchtop with the unaided eye. A laser pointer also produces, very nearly, plane wavefronts, which is one of the reasons it will appear in this chapter so often. The second picture is the **spherical wavefront**: a sphere expanding outward, like the surface of an inflating balloon, centred on a single point source. The headlamp of a daladala bus on Sokoine Drive at midnight produces spherical wavefronts, at least until the light hits a wall and starts to do other things which we will discuss in due course.

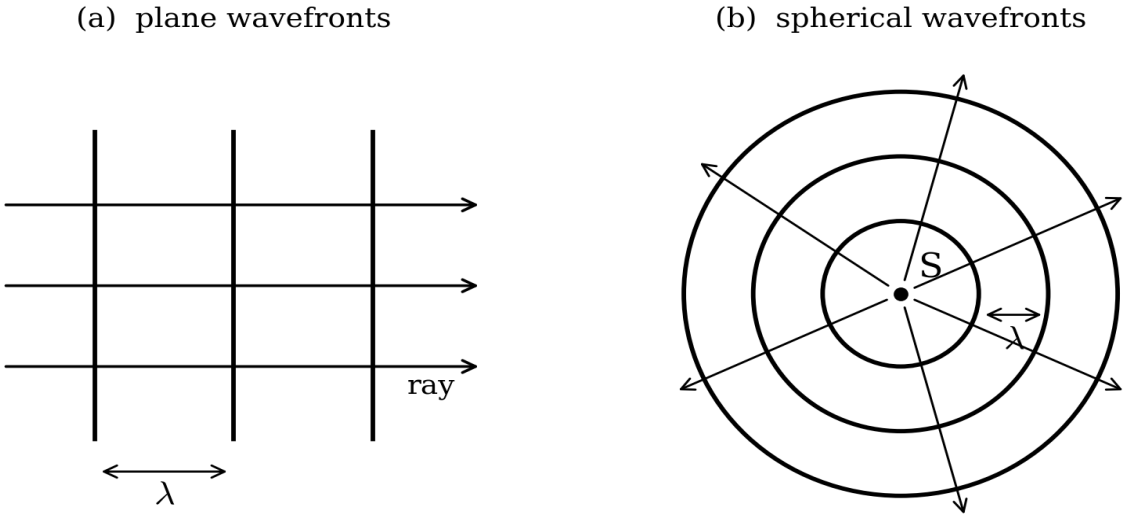


Figure: ***Plane and spherical wavefronts.*** *Wavefronts are drawn as solid lines; rays, the perpendicular bookkeeping arrows, show the direction in which the wave is travelling.*

Notice in the figure above the small arrows that we have drawn perpendicular to the wavefronts. These arrows point in the direction the wave is travelling, and we call them **rays**. A ray is therefore **not** a fundamental physical object; it is just a perpendicular bookkeeping mark on a wavefront, telling us which way the wavefront is advancing. The geometric optics you met in earlier years drew nothing but rays and got remarkably far. Wave optics will draw both wavefronts and rays, and it is the wavefronts that will carry the new physics.

Why did the ray-only picture succeed for so long without ever talking about wavefronts? Because for ordinary optical apparatus the wavelength of light is so much smaller than the apertures and the obstacles around it that nothing interesting happens at the scale of one wavelength. Wavefronts then advance like a tightly packed deck of playing cards, and the only thing worth tracking is which way the deck is moving, namely the ray. Wave optics, as a useful subject, begins exactly where this approximation breaks. We will hit that breaking point twice in this chapter: once when an aperture shrinks to the order of λ, and once again when two waves recombine after a brief separation. Both events demand wavefronts; neither can be solved by rays alone. So we install the wave machinery now and rent it back to ourselves over and over for the rest of the chapter.

The single physical idea that will let us do this is what Christiaan Huygens, a Dutch physicist, proposed in 1678. He suggested that any wavefront could be replaced, with no loss of physical content, by a swarm of tiny secondary sources scattered over its surface. Each secondary source then sends out its own little spherical wavelet, the wavelets travel forward at the speed of the wave, and the next wavefront is whatever shape is tangent to all those wavelets a moment later. A wavefront, in other words, generates the next wavefront by sub-contracting the work to a million small spheres. Here is the principle stated cleanly.

Huygens' principle: *Every point on a wavefront acts as a source of secondary spherical wavelets that travel forward at the same speed as the wave. The next position of the wavefront, after a short time Δt, is the surface tangent to all those wavelets at that later instant.*

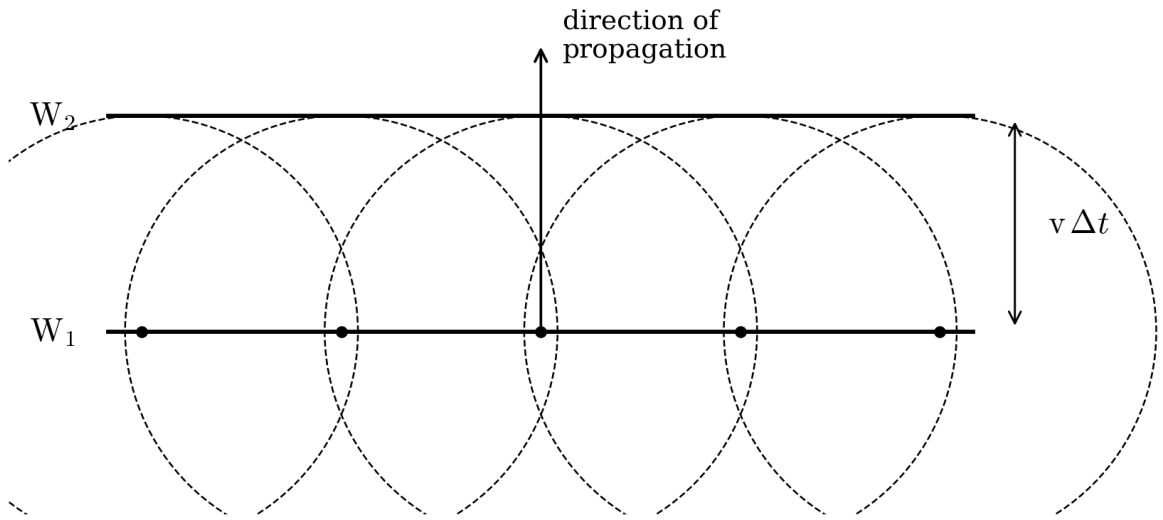


Figure: ***The Huygens construction.*** *Five secondary sources are placed on the wavefront W_1; after a time Δt each has emitted a spherical wavelet of radius $v\,\Delta t$; the new wavefront W_2 is the line tangent to the forward edge of all the wavelets.*

Three observations about the principle deserve a moment before we put it to work.

First, *the secondary wavelets are spherical because the medium around them is the same in every direction;* the wavelet does not know which way the wavefront was originally going, so it spreads in all directions equally. The original direction of the wave is recovered only when the envelope is taken: the envelope on the forward side is the new wavefront, and the would-be envelope on the backward side, mathematically, would correspond to a wavefront moving back into the wave that has already passed by. Huygens did not have an explanation for why we should ignore the backward envelope. The explanation came later, with Maxwell's equations, and amounts to saying that the backward wavelets cancel each other out leaving only the forward envelope. We will accept this on physical grounds for now.

Second, *the principle is* ***scale-free***. It does not care whether the wave is a sound wave in air, a ripple on Lake Victoria, or a beam of red laser light. The same construction handles all of them, with only the speed v changing.

Third, and most importantly for what is to come, *the principle is a* ***prediction machine***. Given the wavefront now, it tells us the wavefront a moment later. It is the wave-mechanical sister of Newton's second law of motion: $F = ma$ told you what happens to a particle next; Huygens tells you what happens to a wave next. From this single statement we will pull out *the law of reflection, Snell's law of refraction, the Young's double-slit fringe pattern, the single-slit diffraction pattern, the diffraction grating equation,* and *the Newton's rings formula.* Six results, one principle, no exceptions. It is the bargain of the chapter, and it costs only one figure on a piece of paper.

Let us spend the principle on something the student already believes, by way of a sanity check: *a flat wavefront travelling through a uniform medium continues forward as a flat wavefront.* This is the simplest possible test, and if Huygens fails it we have no business going any further with the chapter.

Take a horizontal flat wavefront W_1 at time t. Decorate it with five secondary sources, evenly spaced. After a short time Δt, each source has emitted a wavelet of radius $v\Delta t$. The envelope of these five circles, on the forward side, is a horizontal straight line a distance $v\Delta t$ above W_1. That is the new wavefront W_2. It is flat, it is parallel to W_1, and it has advanced by $v\Delta t$. The wave model contains the ray model exactly as the previous section promised. Tick.

Sanity check passed. The principle now needs to earn its keep. We will spend it on two laws the student already knows but has never derived: **the law of reflection** and **Snell's law of refraction**. By the time we finish, both laws will be understood in a way that no amount of memorisation could ever achieve.

Reflection from Huygens

The next figure shows a plane wavefront AB approaching a flat reflecting surface XY at an angle. The wavefront is drawn at the very instant the leading point A first touches the surface. At that same instant, the trailing end of the wavefront, point B, is still in the air, with some travel left to do before it too reaches the surface.

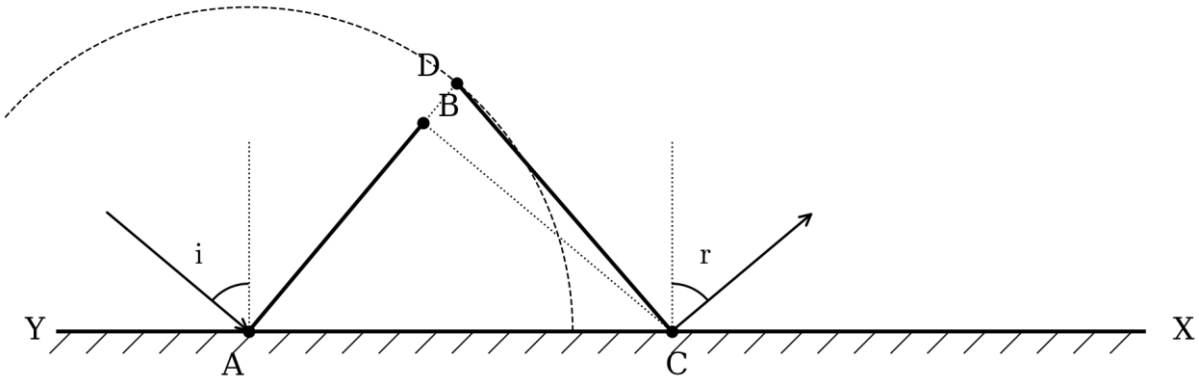


Figure: ***Reflection by the Huygens construction.*** *AB is the incident wavefront at the moment A first touches the surface. The wavelet emitted from A grows to radius BC during the time the trailing edge of the wavefront takes to travel from B to C. The reflected wavefront CD is tangent to the wavelet at D.*

Two physical facts go into the construction, and they are worth saying out loud before we touch any algebra.

The first fact is **timing**. The trailing point B reaches the surface at C only after a time $\Delta t = BC/v$, where v is the speed of the wave in the medium above the surface and BC is the distance from B to C measured along the path the trailing edge of the wavefront has yet to travel. During this same time interval, the wavelet that started at A has been spreading; its radius at the moment B arrives at C is $v\,\Delta t = BC$. So the wavelet at A has radius BC exactly when B has just arrived at C.

The second fact is the **boundary condition**. The reflecting surface is impenetrable; the wavelets cannot propagate into it. The new wavefront, therefore, is the envelope of secondary wavelets on the side of the surface that the original wave came from.

Now the construction. Draw a circle of radius BC centred at A. This is the wavelet that started when A first touched the surface, drawn at the same later instant when B has arrived at C. From C, draw the tangent to that circle. The tangent meets the circle at some point D. The reflected wavefront at this instant is the line CD.

The geometry now does the work for us. Triangle ABC is right-angled at B, with the angle at A equal to the angle of incidence i (the angle between the incident wavefront and the surface, equivalently the angle between the incident ray and the normal). Triangle ACD is right-angled at D, because CD is tangent to the circle and AD is the radius drawn to the tangent point. Both triangles share the hypotenuse AC. Therefore:

$$\mathbf{\sin i} = \frac{BC}{AC} = \frac{AD}{AC} = \mathbf{\sin r}$$

where r is the angle the reflected wavefront CD makes with the surface XY, and we have used $AD = BC$ because both equal the wavelet radius. The two sines are equal, so the two angles are equal. The angle of incidence equals the angle of reflection. Done.

That is the law of reflection, derived in three lines from the principle of a Dutchman who lived three hundred and fifty years ago. Notice what we did and did not need: we needed the wave model to give us the wavelets, we needed a tangent-and-radius right-angle, and we needed two right triangles to share a hypotenuse. We did not need to memorise anything. The phenomenon now belongs to you in a way that *just memorise i = r* never could.

Refraction from Huygens: Snell's law

The same construction repeats for refraction, with one new ingredient: the speed of the wave in the second medium is different from the speed in the first. Call them v_1 and v_2. As we will see in a moment, it is the change in speed that produces the change in direction. Speed is the cause; bending is the symptom.

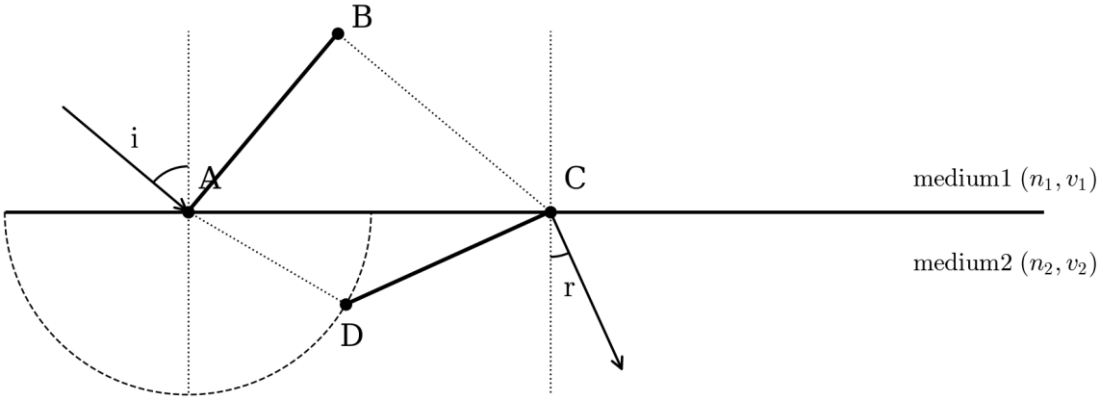


Figure: ***Refraction by the Huygens construction.*** *The wavelet from A propagates into medium 2 with speed v_2, so it has reached only a radius $(v_2/v_1)\cdot BC$ by the time the trailing edge of the wavefront has arrived at C in medium 1. The refracted wavefront CD is tangent to the wavelet at D in medium 2.*

The figure above shows a plane wavefront AB approaching a flat boundary between two media at an angle. As before, A is the leading point of the wavefront that first touches the boundary; B is the trailing point still in medium 1. While B travels the remaining distance BC in medium 1, taking a time $\Delta t = \frac{BC}{v_1}$, the wavelet that started at A is propagating into medium 2 at the new speed v_2. By the time B has arrived at C in medium 1, the wavelet at A has reached the radius AD given by:

$$AD = v_2 \Delta t = \frac{v_2}{v_1} BC$$

Draw, then, a circle of this radius centred at A but on the medium-2 side of the boundary. From C draw the tangent to that circle; the tangent meets the circle at some point D. The refracted wavefront at the instant B arrives at C is the line CD.

Two right-angled triangles, sharing the hypotenuse AC. In triangle ABC, $\sin i = \frac{BC}{AC}$. In triangle ACD, $\sin r = \frac{AD}{AC}$, and $AD = \frac{v_2}{v_1} BC$. Dividing the two sines,

$$\frac{\sin i}{\sin r} = \frac{v_1}{v_2}$$

Two practical reorganisations of this result. First, define the **refractive index** of medium 1 by $n_1 = c/v_1$, and of medium 2 by $n_2 = c/v_2$, where c is the speed of light in vacuum. Then:

$$n_1 \sin i = n_2 \sin r$$

which is Snell's law in the form your final examination will use. Second, observe that in passing from medium 1 to medium 2 the wavelength scales as $\frac{\lambda_2}{\lambda_1} = \frac{v_2}{v_1}$, while the frequency does not change at all. The frequency cannot change at a boundary because oscillations cannot be created or destroyed there, a result we already proved at the end of Chapter 7. The refracted wavefronts in the diagram, drawn one wavelength apart, are spaced more closely if the wave has slowed down ($n_2 > n_1$), and more widely apart if it has sped up. The bend, the speed change, and the wavelength change are not three different things; they are three views of the same one thing.

Two worked examples will now exercise the new machinery, one numerical and one explanatory, while the ideas are still fresh in the head.

BINDER Example 1

Light from a laser pointer enters a basin of clean water from above, making an angle of 50° with the vertical. Calculate the angle the refracted beam makes with the vertical inside the water. The refractive index of water is 1.33.

Solution

Choose the vertical as the normal to the water surface. The angle of incidence in air is $i = 50°$. The angle of refraction in water, the unknown we want, is r. Treat air as vacuum for this purpose, so $n_1 = 1.00$, and $n_2 = 1.33$ for water. Apply Snell's law,

$$n_1 \sin i = n_2 \sin r$$

Rearrange for sin r:

$$\sin r = \frac{n_1}{n_2} \sin i = \frac{1.00}{1.33} \times \sin 50° = (0.7519) \times (0.7660) = 0.5760$$

$$r = \sin^{-1}(0.5760) = 35.2°$$

The refracted beam therefore makes an angle of 35.2° with the vertical inside the water.

Making Sense of the Answer: *The refraction angle (35.2°) is smaller than the incidence angle (50°), so the beam bent towards the normal as it entered the denser medium. That is exactly what the Huygens construction told us to expect: the wavelets in water are slower than the wavelets in air, so the side of the wavefront that hit the water first slowed down first and dragged the rest of the wavefront after it, like one wheel of a daladala bus dropping into a soft patch on the road and pulling the whole bus a little to the side. The same physics, the same answer. The laser pointer, sadly, will not work better in air for having been so wet.*

Think Like a Physicist: *Whenever you compute a refraction angle and the question gives you the refractive indices, write Snell's law in the form $n_1 \sin i = n_2 \sin r$ before you do anything else. Many marks have been lost over the years by students who wrote the inverse form by mistake and then refused to notice that their answer placed the refracted ray bending the wrong way. The cure is to draw a quick sketch with the normal, mark which medium is denser, and check at the end whether the ray bent towards or away from the normal. Bending towards the normal means going into a denser medium. Bending away means going into a less dense one. The arithmetic should agree with the geometry, every single time.*

REAL Example 2

On a clear afternoon on the Dodoma–Singida highway, the tarmac shimmers in the distance like a sheet of standing water. As the bus approaches, the puddle retreats; by the time the bus is over the spot where the puddle was, the spot is dry. Explain, using wavefronts and the dependence of light speed on the medium, why the driver sees water that is not there. Identify what physical quantity varies with height above the road and how that variation produces the illusion.

Solution

The starting point is the layer of air just above the tarmac. Several hours of midday sun have heated the road, the road has heated the air sitting on it, and the temperature of the air decreases smoothly with height for the first one or two metres above the surface. Density of a gas falls as its temperature rises, and a less dense gas has a lower refractive index. The refractive index of air therefore increases smoothly with height above the road. There is no sharp boundary anywhere; the index changes continuously with height.

Now turn to the wavefront. A ray of light coming from the sky and travelling almost parallel to the road carries a wavefront whose lower edge sits in the hotter, lower-index, higher-speed air, and whose upper edge sits in the cooler, higher-index, lower-speed air. By the Huygens construction we have just used twice, the lower edge of the wavefront advances faster than the upper edge during every short interval Δt. The wavefront tilts as it travels: the ray, perpendicular to the wavefront, bends upward, away from the road.

The driver's eye receives this bent ray. The eye, like every eye, assumes that light has travelled in a straight line from its source. It projects the incoming ray backward along that straight line to a virtual source somewhere on the road in front of the bus, and reads off that virtual source as the origin of the light. The virtual source is a piece of the sky, apparently reflected off the tarmac: a sky-blue puddle, complete with shimmering edges where the temperature gradient is unsteady.

Why does the puddle retreat as the bus approaches? The geometry that places the virtual sky-image on the road in front of the bus depends on the angle at which the bent ray reaches the driver's eye, and that angle depends on the driver's distance from the apparent puddle. As the bus moves forward, the geometry shifts; the spot of road on which the virtual image lands shifts forward by the same amount; and the puddle slides forward in step with the bus. There is never any water; only a continuous refractive index pretending to be one.

The physical quantity that varies with height is the refractive index of the air. The variation is smooth, with no sharp boundary. The variation produces the illusion because the lower part of every wavefront crossing the gradient travels faster than the upper part, the wavefronts tilt upward, the rays curl up and away from the road, and the eye, refusing to believe in curving rays, projects them back to a virtual source on the tarmac.

Making Sense of the Answer: *This is the wave-mechanical account of a mirage, and the only physics it spends is what we already had two pages ago: speed depends on the medium, wavefronts move with that speed, and a wavefront whose two ends move at different speeds tilts. The same picture explains why the sun is geometrically below the horizon for nearly half a minute during which we still see it (the atmosphere bends the rays around the curve of the Earth), why stars appear to twinkle (random fluctuations of the index with altitude), and why a graded-index optical fibre can carry your data across a continent (the index is engineered to vary smoothly across the fibre's cross-section, so that all rays arrive in step). Whenever a medium varies smoothly, divide it in your head into infinitely many infinitely thin layers of constant index, apply Snell's law at each interface, and take the limit. This is what calculus looks like in optics.*

Think Like a Physicist: *When you encounter a phenomenon that looks like reflection without a reflector, check first whether the medium itself has a smoothly varying refractive index. Mirages on a hot road, the apparent flattening of the setting sun, the bending of starlight at low altitude, and the guiding of light along a graded-index fibre are all variations on a single theme. The mathematics is always the same: the wavefront is tangent to a swarm of secondary wavelets, the wavelets travel*

faster where the index is lower, and the wavefront curls toward the higher-index region. No new physics; only the same physics, paid out continuously instead of in lumps.

We have now spent the principle three times. Once on a sanity check (a flat wavefront stays flat), once on the law of reflection, and once on Snell's law of refraction. The reward is that none of these laws is a memorised slogan any more. They follow, in three lines each, from the same physical statement that every point on a wavefront is a tiny secondary source. The next partition of this chapter asks the deeper question. *What happens when the wavefront meets an obstacle whose size is comparable with one wavelength?* That is the question Huygens answers most spectacularly, and it is where physical optics finally separates itself from geometric optics for good.

COHERENCE AND THE CONDITIONS FOR INTERFERENCE

Some weeks after the assembly-hall encounter with the CD, Kipute had decided to verify, on his own, the most surprising claim Mr. Akilikubwa had made all term. The teacher had said, repeatedly and with the enthusiasm of a man certain that he was about to be doubted, that two beams of light coming from the same source by two different paths can add up to darkness in some places and to double brightness in others. Kipute did not believe it. So one Saturday afternoon, with the school empty and the laboratory door unlocked, he borrowed two identical school torches, pointed them at a dark wall, and looked very hard for the alternating dark and bright stripes that ought to appear where the two beams overlapped. He saw, instead, a single uniform circle of light, slightly brighter than either torch alone. He looked harder. The circle stayed uniform. He moved the torches closer together, then further apart, dimmed them, brightened them, walked across the room. The wall stayed uniformly lit. After half an hour, he gave up.

Mr. Akilikubwa, who happened to be returning a borrowed laptop to the Headmaster's office, paused at the laboratory door.

Kipute: *"Sir, you said two beams of light could give bright and dark stripes."*

Mr. Akilikubwa: *"I did say that, Kipute. Repeatedly."*

Kipute: *"I cannot make it happen with two torches."*

Mr. Akilikubwa: *"You will not, ever, with two torches. You could spend your entire life pointing torches at walls and you will not see fringes. Not because the physics is wrong, but because the torches are not the right kind of source. The previous chapter laid out the principle waves obey when they meet, which is superposition, but it left out the conditions sources must satisfy for the resulting pattern to be visible to a human eye. The most important of those conditions has a name we have not yet used in the book, and the name is* ***coherence.*** *Two torches, as you have just discovered, are* ***not*** *coherent. Two slits illuminated by the same lamp are coherent. A laser pointer, which we have not yet seen but soon will, is so coherent that two of its echoes off opposite walls of a corridor will give you fringes between themselves on a third wall behind your back. The whole of this section is about explaining what coherence is, why two ordinary lamps do not have it, and what physicists do about it. Bring your torches with you to next Monday's lesson, Kipute. We are going to use them as the negative example."*

Phase difference: A quick recap

Before we install coherence we owe ourselves one minute on what was already proved in the previous chapter. The principle of superposition told us that when two waves arrive at the same point of a medium their displacements add. The relation between the path-difference Δx of the two paths and the phase-difference $\Delta\varphi$ of the two waves at that point is:

$$\Delta\varphi = \frac{2\pi}{\lambda}\Delta x$$

Constructive interference occurs when $\Delta\varphi$ is an even multiple of π (so the path difference is a whole number of wavelengths). Destructive interference occurs when $\Delta\varphi$ is an odd multiple of π (path difference of half-wavelengths). At a fixed point in space, with two fixed paths, the value of $\Delta\varphi$ is fixed by geometry, so each point on a screen has a definite condition: bright, dark, or somewhere between.

Hidden inside that last sentence is one assumption we have been making without saying so out loud. The assumption is that the two waves we are adding *share a stable phase relationship*. If they did not, $\Delta\varphi$ at any fixed point would not be fixed; it would wander; and the bright-dark pattern would slide around faster than the eye could follow. Kipute's two torches are exactly this case. The whole of this section is about why.

Coherence

Two sources are **coherent** *if the phase relationship between the waves they emit stays constant in time, or at least changes only on a timescale much longer than the time over which we are observing.* They are **incoherent** *if the phase relationship wanders unpredictably on a timescale shorter than ours.*

To make this picture concrete, look at the figure below. The top panel shows the phase difference between two coherent sources plotted against time: it is flat, because the two sources march in lock-step and the difference between their phases never changes. The bottom panel shows the phase difference between two incoherent sources plotted against the same axes: it is a step function, jumping at random to a new value every short interval of duration τ_c. That short interval has a name; it is called the **coherence time** of the source.

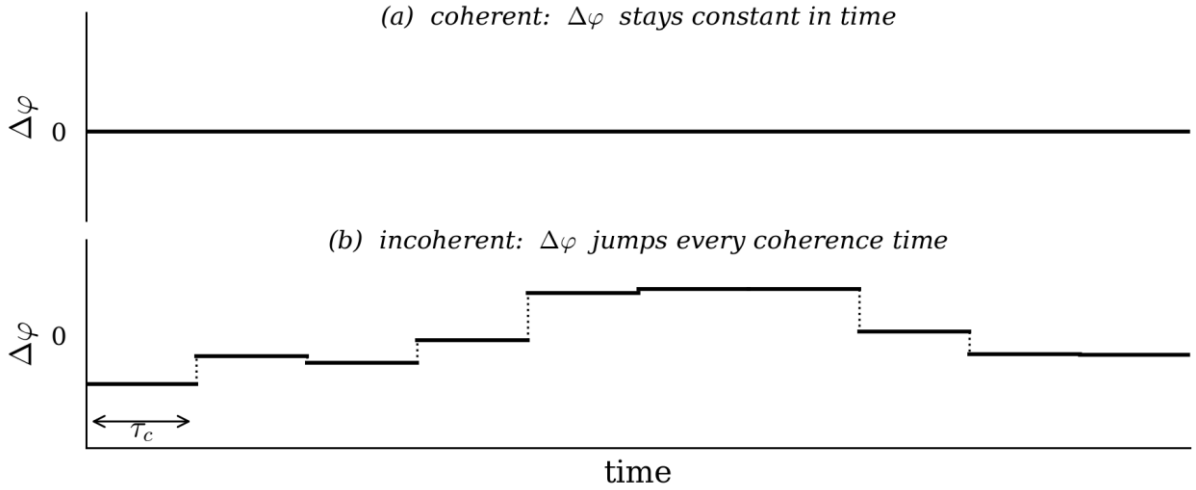


Figure: ***Phase difference between two sources, plotted against time.*** *Coherent sources keep a constant phase difference; incoherent sources reset their phase difference to a random new value every coherence time τ_c.*

How long is the coherence time? It is set by the spread Δf of frequencies the source emits. A source that emits a perfectly single frequency would have an infinite coherence time and a wave that marched in step with itself forever. A real source emits a small range of frequencies; the different frequencies drift out of step with one another, and the coherence time is roughly the time it takes them to drift by one cycle:

$$\boldsymbol{\tau_c \approx \frac{1}{\Delta f}}$$

The narrower the frequency spread, the longer the coherence time. The **coherence length** of the source *is the distance the wave travels in one coherence tim*e,

$$\ell_c = c\,\tau_c \approx \frac{c}{\Delta f} \approx \frac{\lambda^2}{\Delta\lambda}$$

where the last form, equivalent because of $c = f\lambda$, is the version most often quoted in laboratories that measure light by wavelength rather than by frequency. Three real-world numbers fix the orders of magnitude. A sodium discharge lamp at room temperature has a frequency spread of about 1.5×10^9Hz, giving a coherence length of about 20cm. A white-light incandescent bulb spans the whole visible spectrum, Δλ of about 300nm, giving a coherence length of about 1μm. A common helium-neon laser has a frequency spread of a few hundred kHz and a coherence length of several hundred metres. The factor of 30 million between the bulb and the laser is what every interference experiment in this chapter has to negotiate.

Why two independent lamps do not interfere

Now we can resolve Kipute's puzzle. A torch is a hot-filament lamp; the atoms in the filament emit light independently, and any single emission lasts only as long as the coherence time, after which the next emission begins from a fresh, randomly-chosen phase. The coherence time of a torch is comparable with that of a tungsten incandescent bulb, of order 10^{-14}s. Two torches, even if they emit at exactly the same wavelength, randomise their phases independently of each other. Their relative phase, the quantity Δφ the human eye would need to sit still in order to see fringes, jumps to a new random value every femtosecond.

The eye and the brain together integrate the light reaching them over a period of about 50 milliseconds, the so-called persistence of vision. During those 50 milliseconds, the two torches' relative phase has reset itself roughly 5 trillion times. The fringe pattern shifts trillions of times during one moment of seeing. The eye averages over the trillions of shifted patterns, and the average is uniform brightness. Kipute's wall, accordingly, looked uniform. There were fringes there, but they were sliding past too fast for the eye, the camera, the photographic plate, or any benchtop detector to catch.

The condition for fringes to be visible, then, is that *the observation time of the detector should be less than the coherence time of the source*. Either the detector must be unimaginably fast (faster than 10^{-14}s for two torches), or the source's coherence time must be made enormous, or, by far the cleanest workaround, the two waves must be made to share a single source's phase by construction. The third route is the one all the experiments in this chapter take, and it has its own name.

The split-source trick

Take a single source, monochromatic and bright. Split its light into two paths. Recombine the two paths at the screen. Whatever phase the source happens to be in at any one instant, both paths inherit that phase, and the relative phase between them at the screen depends only on the geometric difference of their path lengths. The source's phase wandering, however violent, cancels out: it is added to both waves and subtracted in the difference. The two paths, by construction, are perfectly coherent with each other for all time, regardless of the coherence time of the source itself.

There are exactly two ways to perform the split, and they cover every interference experiment we will meet for the rest of this chapter and every interference experiment that has ever been performed. The following figure shows them side by side.

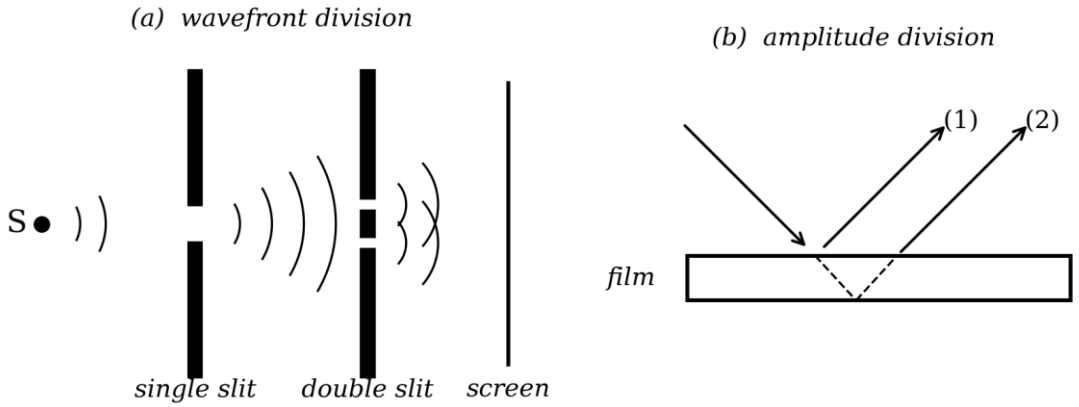


Figure: ***The split-source trick.*** *(a) Wavefront division: a single wavefront is sliced into two by passing it through two narrow slits. The slits act as new coherent sources. (b) Amplitude division: a single ray is partially reflected at one surface and partially at another, so the same wave returns by two paths.*

In **wavefront division**, panel (a) of the figure, the source S sends light through a single narrow slit, which guarantees that the light arriving at the next barrier is in identical phase across its width. That second barrier carries two narrow slits, S_1 and S_2, which act as two new coherent sources. The geometry of the screen behind them gives different path lengths for waves arriving from S_1 and S_2, and the path-difference logic produces the bright-dark pattern. Young's double-slit experiment, the Fresnel biprism, the Fresnel double mirror, and Lloyd's single mirror all belong to this family. We will treat Young's apparatus in detail in the next section.

In **amplitude division**, panel (b) of the figure, a single ray is split into two by partial reflection at the surfaces of a thin film. One reflection comes off the top surface; the other reflection comes off the bottom surface, after a round-trip through the film of thickness t and refractive index μ. The optical path-difference between the two reflections is $2\mu t$ plus a small contribution from any phase change at the surfaces (a contribution we will earn properly in a later section). Soap bubbles, oil films on rainwater, anti-reflective coatings on a camera lens, and the Newton's-rings apparatus all belong to this family. We will treat the latter two in their own sections.

The four conditions for sustained interference

We can now state, in clean numbered form, the conditions an experimental setup must satisfy for a stable bright-dark pattern of fringes to appear. There are four of them; the first one is the new physics of this section, the other three follow from common sense once it is on the page.

First condition: *The two waves must be* ***coherent****.* This means in practice that they must come from the same source by way of one of the two splitting tricks above. Two independent sources of any kind, even of identical wavelength, will not do.

Second condition: *The two waves should be* ***monochromatic****,* or near enough to it. A wide spread of wavelengths gives each colour its own fringe spacing, and the colours overwrite one another into a uniform glow after only a few orders. White light gives a few coloured fringes very close to the central bright fringe and washes out beyond. Sodium light gives several hundred clean fringes. A laser gives thousands.

Third condition: *The two waves should have* ***comparable amplitudes****.* The minima of an interference pattern are dark only when destructive interference is complete, and complete cancellation requires equal amplitudes. If one wave is much weaker than the other, the minima are merely dim, the contrast is poor, and the fringes are difficult to see even when they are there.

Fourth condition: *The two waves should have the* ***same polarisation state****.* Two waves whose electric fields oscillate in mutually perpendicular planes do not interfere with each other; they superpose into a wave with a tilted polarisation, but the intensities at every point on the screen simply add as for incoherent waves. We will return to polarisation in its own section later in this chapter; for now, accept that the experiments we are about to describe all use light whose polarisation is the same on both paths, because the same source feeds both paths and partial reflection does not turn one polarisation into the other.

Two worked examples now follow, because ideas of this kind are understood best, and enjoyed best, when they are put to work.

BINDER Example 3

A sodium discharge lamp emits its yellow line at $\lambda = 589\text{nm}$ with a frequency spread $\Delta f \approx 1.5 \times 10^9\text{Hz}$ arising mostly from Doppler broadening at operating temperature. Calculate (a) the coherence time and (b) the coherence length of this

source. Comment on whether the lamp is suitable for a Young's double-slit experiment in which the path-difference between the two slits and the screen never exceeds a few millimetres.

Solution

(a) The coherence time follows directly from the frequency spread:

$$\tau_c \approx \frac{1}{\Delta f} = \frac{1}{1.5 \times 10^9\ \text{Hz}} = 6.7 \times 10^{-10}\ \text{s} = 0.67\text{ns}$$

(b) The coherence length follows from the coherence time and the speed of light:

$$\ell_c = c\,\tau_c = (3.0 \times 10^8 \text{m/s})(6.7 \times 10^{-10}\text{s}) \approx 0.20\text{m} = 20\text{cm}$$

The coherence length is 20cm.

Path differences of a few millimetres in a Young's apparatus are smaller by a factor of about a hundred, so the lamp is comfortably suitable. We expect, and we will see, hundreds of clean fringes from a sodium lamp behind a Young's double slit.

Making Sense of the Answer: *A coherence length of 20cm is a good benchtop number. It tells us that fringes will survive any path difference up to about 20cm before they wash out, which leaves enormous room for the typical Young's apparatus geometry of slit-separation under one millimetre and screen distance under one metre. It also tells us that the same lamp would not work for a Michelson interferometer with arms of a few metres, or for a holographic exposure that needs path-matching to better than a fraction of a wavelength: those experiments demand laser sources, whose coherence lengths are measured in hundreds of metres rather than centimetres. The single quantity ℓ_c tells us, before we do any geometry, what kind of interference experiment a given source can support.*

Think Like a Physicist: *When you are handed an interference problem, the first quantity to compute is the coherence length of the source. Compare it with the geometric path differences of the apparatus. If the path difference is much less than ℓ_c, the experiment will work and you can reach for the geometry. If the path difference is comparable with ℓ_c, expect a few faint fringes only. If the path difference exceeds ℓ_c, expect no fringes at all and look for a different source. The coherence length is the licence the source carries; the geometry is what you do with the licence.*

REAL Example 4

A soap film of thickness about 500nm and refractive index 1.33, illuminated by ordinary daylight, displays bright iridescent colours that shift as the film thins. A clean window pane of thickness 3mm and refractive index 1.50, illuminated by exactly the same daylight, displays nothing of the kind. Both surfaces have two reflecting interfaces (front and back), so both ought, on the face of it, to give amplitude-division interference. Explain, using the coherence length of daylight, why one shows colours and the other does not.

Solution

Estimate first the coherence length of daylight. Visible light spans the wavelength range from about 400nm to 700nm, so the spread is $\Delta\lambda \approx 300\text{nm}$ centred on $\lambda \approx 550\text{nm}$. The coherence length is:

$$\ell_c \approx \frac{\lambda^2}{\Delta\lambda} = \frac{(550\text{nm})^2}{300\text{nm}} \approx 1.0 \times 10^{-6}\text{m} = 1\mu\text{m}$$

Daylight has a coherence length of about one micrometre. Now compare it with the optical path differences in the two films.

For the soap film, the optical path between the front-surface reflection and the back-surface reflection is $2\mu t = 2(1.33)(500\text{nm}) \approx 1.3\mu\text{m}$. This is just barely larger than the coherence length, so the two reflections from the soap film still share a partial phase relationship. The interference is real but weak; only a few orders survive. Those few orders, with the path difference depending on the thickness t of the film and on the wavelength, give different colours at different thicknesses. As the film drains under gravity, the thickness varies smoothly with height and the colours move with it. We see iridescent bands.

For the window pane, the optical path is $2\mu t = 2(1.50)(3\text{mm}) = 9\text{mm} = 9\times10^3\mu\text{m}$. This is nine thousand times larger than the coherence length of daylight. By the time the back-surface reflection emerges from the glass, it has lost any memory of the phase the front-surface reflection had when it began. The two reflections are now mutually incoherent, even though they came from the same source, because they have spent more than a coherence time apart. Their intensities simply add. There is no interference and no colour, only ordinary glass with two reflective surfaces.

Making Sense of the Answer: *This is why thin films display colour and thick plates do not. The threshold thickness is set by the source's coherence length divided by twice the refractive index, which for daylight on glass is about 0.3μm. Any film thinner than this gives interference; any film thicker than a few microns does not. Soap bubbles fall right at the limit, which is why they are vivid and ephemeral; the iridescent shimmer on a petrol-rainbow puddle the morning after rain works the same way; the anti-reflection coating on a camera lens, designed to be a quarter of one wavelength thick, sits well inside the limit and exploits it. The same physics, the same threshold, three different applications.*

Think Like a Physicist: *When you see two reflecting surfaces and you are wondering whether they will give interference, do not ask first whether the two waves are coherent at the source. Ask whether the optical path difference between the two*

reflections is smaller than the coherence length of the source. If yes, expect interference. If no, expect ordinary reflections that add in intensity, not in amplitude. The geometric thickness of the film is the diagnostic; everything else follows.

We have now installed the new physics this chapter needed: coherence, the split-source trick, and the four conditions a real apparatus must satisfy for a stable fringe pattern to be visible. The next section picks up from here and treats, in geometric detail, the most famous and most heavily examined case of wavefront division: Thomas Young's double-slit experiment of 1801. We will derive the fringe-position formula from Pythagoras, work through five different parameter changes and what they each do to the fringe spacing, and finish with the only formula in the chapter that needs to be memorised verbatim because every paper sits one or two of its variants.

YOUNG'S DOUBLE-SLIT EXPERIMENT

Monday morning. Kipute has brought his two torches as instructed, slightly embarrassed at the thought that they are about to be told off in front of the whole class. The physics laboratory has been darkened with thick black cloth over the windows. A sodium lamp sits on a workbench, glowing the unmistakable yellow of every street lamp on Bagamoyo Road; in front of the lamp Mr. Akilikubwa has placed a vertical card with a single narrow slit cut into it; behind the first card and a metre further back stands a second card with two narrow slits scored a fraction of a millimetre apart; behind the second card stretches the length of the laboratory; behind the laboratory stands the white wall opposite the windows. Kipute's torches sit on the desk, switched off, awaiting their negative-example duty.

Mr. Akilikubwa switches on the sodium lamp, turns off the room lights, and asks the class to look at the wall. After a few seconds for the eyes to adjust, the wall is no longer dark. It carries a pattern of faint, evenly spaced, alternating bright and dark stripes, all yellow, running vertically. Kipute, who had spent half his Saturday looking for those stripes with two torches and finding nothing, leans forward and stares.

What the class is looking at is two hundred and twenty-five years old. Thomas Young, the English physician and physicist, set up exactly this apparatus in 1801, with a candle in place of the sodium lamp and a pair of pinholes in place of slits, and used the pattern on the wall to settle, once and for all, an argument that had run for a century and a half: light is a wave, the same kind of wave we have now spent two chapters deriving, and the bright-dark stripes on the wall are the constructive and destructive maxima of two coherent wavefronts adding their amplitudes at every point on the screen. The remainder of this section earns the formula that says exactly where each stripe falls.

The apparatus

The following figure shows the standard arrangement. A monochromatic source S sits on the left. Its light passes through a single narrow slit S_0, which acts as the new effective source for the rest of the experiment. After S_0, the light reaches a second card carrying two narrow slits S_1 and S_2, separated by a small distance d. By the split-source argument of the previous section, S_1 and S_2 act as two new coherent sources, because both inherit their phase from the same wavefront that passed through S_0. The two waves then propagate to a screen placed at a perpendicular distance D behind the double slit. The pattern of bright and dark stripes on the screen is the Young's interference pattern, and every formula in this section is about predicting where on the screen the bright stripes fall.

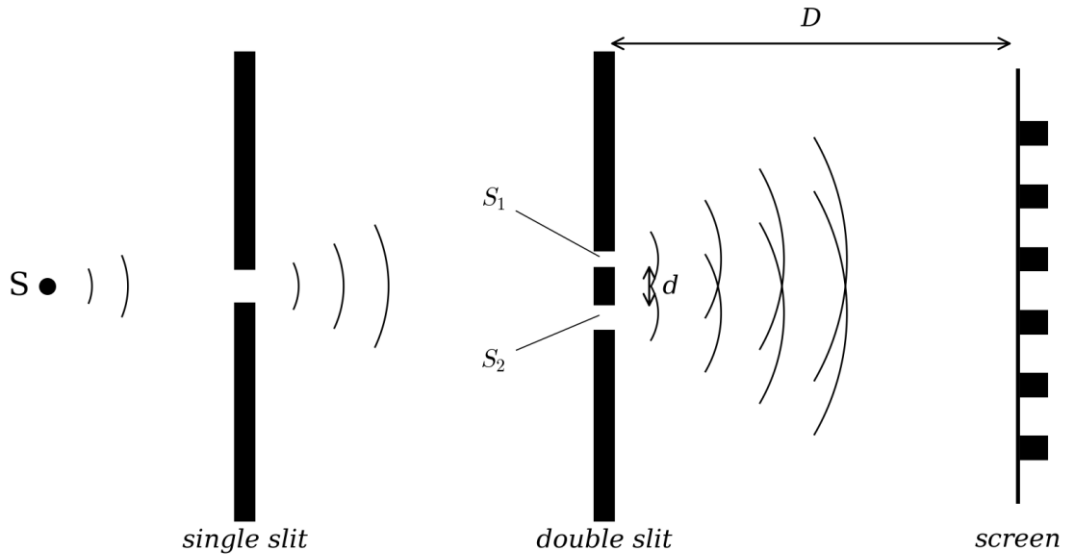


Figure: ***Young's double-slit apparatus.*** *A monochromatic source S sends light through a single slit S_0, which then illuminates a second card carrying two slits S_1 and S_2 separated by d. The two slits act as coherent sources; the screen at distance D behind the double slit displays the interference pattern.*

Three remarks fix the orders of magnitude in your head before any algebra. **First,** *the slit separation d is small*. In a typical school benchtop apparatus d lies between 0.1mm and 1mm; if d is much smaller the fringes spread out so widely that only one or two fit on a sensible-sized screen, and if d is much larger the fringes crowd so close that the eye cannot resolve them. **Second,** *the slit-to-screen distance D is large*. Typical values run from 50cm to a few metres; D much larger than d is the geometric assumption that allows the small-angle approximation we are about to make. **Third,** *the wavelength of light is tiny*. For the sodium lamp on the bench, $\lambda \approx 589$nm; for a helium-neon laser, $\lambda \approx 633$nm; both numbers are tiny compared with even the smallest slit separation. The chapter's central inequality, ***λ much less than d***, is what keeps the pattern visible at all.

Derivation of the fringe positions

To find where the bright fringes fall we follow the wave from each slit to a general point P on the screen and ask for the path-difference between the two waves there. The geometry is shown below.

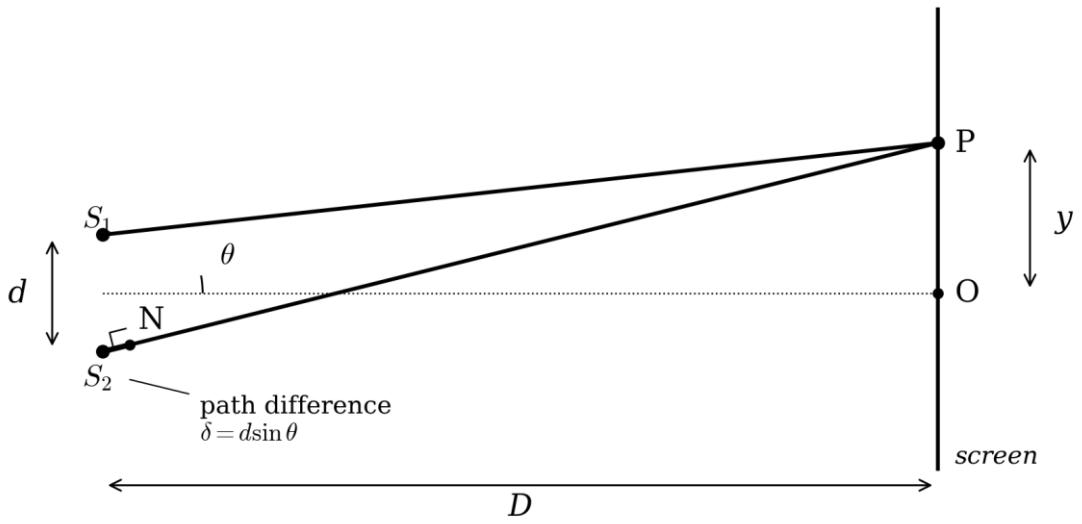


Figure: ***The path-difference geometry.*** *P is a point on the screen at vertical distance y above the central axis O. The path difference $\delta = S_2P - S_1P$ is the segment from S_2 to the foot of perpendicular N dropped from S_1 onto S_2P; in the small-angle limit $\delta \approx d \sin\theta \approx dy/D$.*

Place S_1 and S_2 symmetrically about the central axis with separation d. Place the screen at distance D behind them. Let O be the foot of the perpendicular from the midpoint of the slits onto the screen, and let P be a general point on the screen at vertical distance y above O. We want the path-difference,

$$\delta = S_2P - S_1P$$

Drop a perpendicular from S_1 onto the line S_2P, and call the foot of that perpendicular N. In the far-field limit, where the screen is much further than the slit separation, the two paths S_1P and NP are very nearly equal in length, so:

$$\delta \approx S_2N$$

That is the path-difference, dressed in geometry. In the small-angle limit, where D is much greater than d and the angle θ between the central axis and the line from the midpoint of the slits to P is small, the segment S_2N equals $d\sin\theta$ by the right triangle S_1S_2N. Therefore:

$$\delta = d\sin\theta$$

In the same small-angle limit, $\sin\theta \approx \tan\theta = \frac{y}{D}$, so the path-difference at the point y on the screen is:

$$\boldsymbol{\delta = \frac{d\,y}{D}}$$

This is the master formula of Young's experiment. Every fringe condition follows from it by inserting the right value of δ. Constructive interference requires path-difference equal to an integer number of wavelengths,

$$\delta = n\lambda \quad (n = 0, \pm1, \pm2, \ldots)$$

Combining the two equations and solving for y, the position of the n-th bright fringe is:

$$\boldsymbol{y_n = \frac{n\,\lambda\,D}{d}}$$

The dark fringes lie halfway between the bright ones, with path-difference equal to a half-integer number of wavelengths,

$$\delta = (n + ½)\,\lambda \quad (n = 0, \pm1, \pm2, \ldots)$$

which gives the position of the n-th dark fringe as:

$$\boldsymbol{y_n = \frac{(n + ½)\,\lambda\,D}{d}}$$

These three formulas are the entire output of the section so far. Every numerical Young's-slit problem you will ever solve plugs into one of them. The next subsection notices a quantity that is even more useful in practice than the absolute fringe position, and gives it a name.

The fringe width

The **fringe width**, written **β**, *is the perpendicular distance between any two consecutive bright fringes (or between any two consecutive dark fringes).* From the position formula,

$$\boldsymbol{\beta} = y_{n+1} - y_n = \frac{(n+1)\,\lambda\, D}{d} - \frac{n\,\lambda\, D}{d} = \boldsymbol{\frac{\lambda D}{d}}$$

Three features of this single line are worth pinning down.

First, *the fringe width depends only on the wavelength, the slit-to-screen distance, and the slit separation*. It does not depend on which fringe you happen to be measuring; the bright fringes are evenly spaced and the dark fringes are evenly spaced. The pattern is a clean periodic stripe pattern with one period of width β.

Second, *β scales linearly with λ and with D, and inversely with d*. Doubling the wavelength doubles the fringe width. Doubling the screen distance doubles the fringe width. Doubling the slit separation halves it. These three scaling rules are the most useful single piece of arithmetic in the chapter.

Third, *every fringe width formula* in this chapter, no matter how unfamiliar the apparatus looks, *comes back to β = λD/d in the limit where the screen is far compared with the slit separation.* Newton's rings, single-slit diffraction patterns, diffraction-grating spectra, and the Fresnel biprism all reduce to scaled-up or scaled-down versions of this same formula in the small-angle limit. We are unlocking, in this section, a single physical idea that is going to recur three more times before the chapter is finished.

What happens when you change a parameter

Examiners love the same question dressed up five ways. The question is always: *what happens to the fringe pattern when one parameter of the apparatus is altered?* Five common alterations and their consequences follow.

Change one: *the entire apparatus is **immersed in water** of refractive index n = 1.33.* The wavelength of light in water is λ/n, so the fringe width becomes:

$$\beta' = \frac{(\lambda/n)\,D}{d} = \frac{\beta}{n}$$

The fringes squeeze together by a factor of n. For water this means the new spacing is 1/1.33 ≈ 0.75 of the old spacing. The pattern survives, only tighter.

Change two: *the **slit separation d is reduced***. Since β is inversely proportional to d, halving d doubles the fringe width. Closer slits give wider fringes.

Change three: *the **source slit** S_0 is moved closer to the double slit S_1S_2.* The fringe width β itself is unchanged, because none of λ, D, or d has changed. What changes is the spatial coherence of the wavefront illuminating S_1S_2: as S_0 approaches the double slit, the wavefront curvature across the gap becomes detectable, and the fringes become less crisp. The pattern is the same width but the contrast falls.

Change four: *the **screen is moved closer** to the double slit, halving D*. This is the change that confused Kipanga and Kipute when Mr. Akilikubwa tested them on it during the demonstration.

Mr. Akilikubwa: *"I am about to bring the screen halfway closer. Predict what happens to the fringe width before I do."*

Kipanga: *"Halving the distance to the screen doubles the fringe width. Closer means the fringes spread out."*

Kipute: *"Halving the distance halves the fringe width. The closer the screen, the smaller the spacing."*

Mr. Akilikubwa: *"Two opposite predictions, two confident faces, one formula on the board. Read it. β = λD/d. The fringe width is proportional to D. Half D, half β. Kipute is right and Kipanga is not. The intuition that betrayed Kipanga is the one that says "closer always means more spread out"; that intuition would be correct if the angular spread of the rays was fixed and we were tracking the size of a divergent beam. But the angular spread here is set by sin θ ≈ λ/d, which does not depend on D at all. The spacing of the fringes on the screen is the angular spread times the distance from the source to the screen, so spacing scales with D. Move the screen closer, the angular fan stays the same, the screen samples a smaller piece of it, and the spacing between fringes is correspondingly smaller. Move it further away and the spacing grows."*

He moved the screen halfway closer. The fringes halved in spacing, exactly as the formula said they would.

Change five, and the most important one for examination purposes, is *the introduction of a **thin transparent plate** of refractive index μ and thickness t in front of one of the two slits*. This is the topic of the next subsection because it has its own clean formula and its own family of problems.

The fringe-shift formula

Place a thin transparent plate of refractive index μ and thickness t directly in front of slit S_1, with all other geometry unchanged. The light passing through S_1 now traverses the plate and emerges with an extra optical path of μt instead of the simple geometric path t it would have taken in air. The increase in optical path length, compared with the no-plate case, is:

$$\Delta p = \mu t - t = (\mu - 1)\, t$$

This extra path is added to the wave from S_1 before recombination. The whole fringe pattern therefore shifts on the screen so that the central bright fringe (formerly at y = 0) moves to a new position where the geometric path-difference exactly cancels the extra optical path. From $\delta = dy/D = (\mu - 1)t$, the displacement y:

$$y = \frac{D\,(\mu - 1)\, t}{d}$$

This displacement is called the **fringe shift**, so we rename it:

$$\boldsymbol{\Delta x = \frac{D\,(\mu - 1)\, t}{d}}$$

Three consequences are worth noting. **First,** *the entire pattern shifts as a rigid block;* the fringe spacing β does not change, because the plate adds the same extra optical path to every wave from S_1 regardless of the angle. **Second,** *the shift is towards the side where the plate was placed (*because the central fringe now occurs where the plate-side wave is delayed enough to match the unplated wave). **Third,** *the number of fringes that have crossed any fixed point on the screen during the placement is* $\frac{\Delta x}{\beta} = \frac{(\mu - 1)\, t}{\lambda}$, which gives an experimental method for measuring μ or t by counting the fringes that pass.

With the ideas now simmering nicely, next five worked examples are the moment we lift the lid and finally taste them, where the flavour of the physics rewards the patience.

BINDER Example 5

In a Young's experiment, a helium-neon laser of wavelength 632.8nm illuminates a card scored with two narrow parallel slits 0.20mm apart. The resulting interference pattern is observed on a white screen placed 1.50m behind the slits. Calculate (a) the perpendicular distance from the central bright fringe to the third bright fringe above it on the screen, and (b) the spacing between consecutive bright fringes in the pattern.

Solution

Begin by naming the quantities. Let $\lambda = 632.8\text{nm}$ be the wavelength of the light, $d = 0.20\text{mm}$ the separation between the two slits, and $D = 1.50\text{m}$ the perpendicular distance from the slits to the screen.

(a) The position of the n-th bright fringe is given by:

$$y_n = \frac{n\,\lambda\, D}{d}$$

Substitute n = 3, with all quantities in SI units:

$$y_3 = \frac{3 \times (632.8 \times 10^{-9}\text{m}) \times 1.50\text{m}}{0.20 \times 10^{-3}\text{m}} = 1.42 \times 10^{-2}\text{m} = 14.2\text{mm}$$

The third bright fringe lies 14.2mm above the central maximum.

(b) The fringe width is

$$\beta = \frac{\lambda\, D}{d} = \frac{(632.8 \times 10^{-9}\text{m}) \times 1.50\text{m}}{0.20 \times 10^{-3}\text{m}} \approx 4.7 \times 10^{-3}\text{m} = 4.7\text{mm}$$

The fringes are 4.7mm apart on the screen.

Making Sense of the Answer: *Two checks confirm the answer. The fringe width times three equals the position of the third bright fringe: 3 × 4.7mm = 14.1mm, which matches y_3 = 14.2mm to round-off. The size is also reasonable: a few millimetres is exactly what the eye can comfortably see at arm's length, which is why every Young's apparatus that has ever been built ends up with d in the tenths-of-a-millimetre range and D somewhere between half a metre and a couple of metres. Smaller d gives wider fringes; bigger D gives wider fringes; both choices are made to push the fringes into the comfortable arm's-length range. The apparatus is a piece of pre-tuned arithmetic that someone solved for you, decades before you arrived in the laboratory; it would be ungrateful not to use it. If your computed y_3 comes out in metres rather than millimetres, the school's screen has somehow grown to the size of a football pitch and the calculation needs a unit conversion before anything else.*

Think Like a Physicist: *Always check that the position of the n-th fringe equals n times the fringe width before you trust your numerical answer. The two formulas are derived from the same logic and any rounding error that breaks one will break the other in the same direction. If they disagree by more than the rounding tolerance you have, with overwhelming probability, made an arithmetic slip somewhere in y_3; redo the substitution.*

BINDER Example 6

In a Young's apparatus, a green light source of wavelength 500nm illuminates a card carrying two narrow parallel slits. A travelling microscope is focused on the screen, which sits 1.0m behind the double slit, and counts 21 consecutive bright fringes occupying a horizontal distance of 10mm. Calculate the separation between the two slits.

Solution

Let $\lambda = 500\text{nm}$ denote the wavelength of the light and $D = 1.0\text{m}$ the slit-to-screen distance. The slit separation, which we wish to find, is d.

Twenty-one bright fringes have **twenty** intervals between them (the first and the last fringes bracket twenty fringe-widths). The fringe width is therefore:

$$\beta = \frac{10.0\text{mm}}{20} = 0.50\text{mm}$$

Rearrange the fringe-width formula for the slit separation:

$$d = \frac{\lambda D}{\beta} = \frac{(500 \times 10^{-9}\text{m}) \times 1.0\text{m}}{0.50 \times 10^{-3}\text{m}} = 1.0 \times 10^{-3}\text{m} = 1.0\text{mm}$$

The slits are 1.0mm apart.

Making Sense of the Answer: *There is exactly one trap in this question, and the trap has caught more students over the years than any other detail in Young's slits arithmetic. Twenty-one bright fringes are not twenty-one intervals; they are twenty intervals, the same way twenty-one fence posts make a fence with only twenty panels. Off-by-one in fringe counting halves the fringe width and doubles the slit separation, neatly converting an honest 1.0mm slit separation into a non-existent 0.50mm one. Whenever you read "N bright fringes spanning x millimetres", divide by N − 1, not by N. The fence post is the metaphor that has saved a thousand examination scripts.*

Think Like a Physicist: *Whenever a Young's-slits problem gives you a fringe count instead of a fringe width, your first arithmetic step is to translate count into spacing using "intervals = fringes − 1". Only after this conversion does the formula $\beta = \lambda D/d$ earn its keep. The translation is the kind of small, quiet step that separates a student who has merely memorised the formula from a student who has understood what the formula is counting.*

REAL Example 7

A Young's apparatus illuminated by a sodium lamp produces bright fringes spaced exactly 1.0mm apart on a screen placed 1.0m behind the double slit. The light from the sodium lamp has a wavelength of 589nm. Without disturbing any of the geometry, the entire apparatus is then submerged in a tank of clean water whose refractive index is 1.33. Predict the new fringe spacing, and explain in physical terms why the change occurs.

Solution.

Let $\lambda = 589\text{nm}$ be the wavelength of the light in air, $\beta = 1.0\text{mm}$ the original fringe spacing, $D = 1.0\text{m}$ the slit-to-screen distance, and $n = 1.33$ the refractive index of the water.

In water of refractive index n, the wavelength of light is shorter:

$$\lambda' = \frac{\lambda}{n}$$

The fringe spacing depends on the wavelength of light in the medium where the interference takes place, not on the vacuum wavelength. With d and D both unchanged (these are geometric distances measured in physical space and are not altered by immersion), the new fringe width is:

$$\beta' = \frac{\lambda' D}{d} = \frac{\lambda D}{nd} = \frac{\beta}{n} = \frac{1.0\text{mm}}{1.33} \approx 0.75\text{mm}$$

The new fringe spacing is about 0.75mm. The fringes squeeze together by a factor of $n = 1.33$.

Making Sense of the Answer: *The picture is straightforward once you see what is changing. The geometric path-difference between the two slits and a point at height y on the screen is still dy/D, exactly as it was in air, because none of d, D, or y is altered by the surrounding medium. What is altered is the wavelength used to convert that geometric path-difference into a phase-difference. In air, a path-difference of one wavelength gave the first bright fringe at $y = \lambda D/d$. In water, the wavelength is shorter, so the same first-bright-fringe condition is met at a smaller y, namely $\lambda D/(nd) = \beta/n$. The fringes are tighter because the wave has been compressed by the medium, not because the geometry has changed. The same physics, the same screen, just water in the way of the light.*

Think Like a Physicist: *Whenever a Young's apparatus is altered, ask three questions in this order.* ***First, has the wavelength of light in the medium changed?*** *If yes, β scales as the new wavelength.* ***Second, has the slit separation d changed?*** *If yes, β scales inversely with d.* ***Third, has the screen distance D changed?*** *If yes, β scales linearly with D. The fringe width formula $\beta = \lambda D/d$ is a permission slip with three filled-in fields; whenever any field is altered, β changes in the obvious proportionate way and nothing else is needed. Most exam questions on Young's slits are exactly this template dressed in different clothing.*

HOT Example 8

In a Young's experiment, monochromatic light of wavelength 600nm produces bright fringes that are 1.2mm apart on a screen at some perpendicular distance behind the double slit. Without disturbing the slit separation, the screen is then moved halfway closer to the slits, and the original light source is replaced with a different monochromatic source of another colour. The new bright fringes are observed to be 0.40mm apart on the screen. Calculate the wavelength of the new light.

Solution

Let $\lambda_1 = 600\text{nm}$ be the wavelength of the original source, $\beta_1 = 1.2\text{mm}$ the original fringe width, λ_2 the wavelength of the new source (the unknown we want), and $\beta_2 = 0.40\text{mm}$ the new fringe width. Let D denote the original slit-to-screen distance, so that the new distance is D/2. The slit separation d is the same in both situations, so the ratio $\beta/(\lambda D) = 1/d$ is constant. Apply this twice and divide:

$$\frac{\beta_1}{\lambda_1 \, D} = \frac{\beta_2}{\lambda_2 \, (D/2)}$$

Solve for λ_2:

$$\lambda_2 = \lambda_1 \times \frac{\beta_2}{\beta_1} \times 2 = 600\text{nm} \times \frac{0.40}{1.2} \times 2 = 400\text{nm}$$

The new wavelength is 400nm, which is the deep-violet end of the visible spectrum, very nearly ultraviolet.

Making Sense of the Answer: *Two changes acted on the fringe width here, and both shrank it: D was halved (which on its own would halve β), and the wavelength was shortened (which on its own would shrink β in proportion). The combined shrinkage is exactly 1.2mm to 0.40mm, a factor of three. Halving D gave one factor of two; the rest, a factor of 2/3 in the same direction, came from the colour change. The arithmetic is two steps stacked on top of each other, and the trick is to keep them straight in the head: each parameter pulls β in its own direction, and the final β is the product of the pulls.*

Think Like a Physicist: *When a problem changes more than one parameter at once, never substitute numbers into β = λD/d on each side and divide. Instead, write the formula symbolically for both situations, divide one by the other, and watch the unchanged quantities (here, the slit separation d) cancel out. The cancellation is what kept this problem solvable without knowing d at all. Whenever you find yourself needing a quantity the problem refused to give you, suspect that it is meant to cancel.*

HOT Example 9

A thin glass plate, of refractive index 1.50 and thickness 6.0μm, is placed directly in front of one of the two slits of a Young's apparatus. The two slits are 0.50mm apart, and the screen sits 1.0m behind the double slit. The apparatus is illuminated by yellow light of wavelength 500nm. Calculate (a) the lateral shift of the central bright fringe on the screen, and (b) the number of bright fringes that cross any fixed point on the screen as the plate is being slid into position.

Solution

Let $\mu = 1.50$ denote the refractive index of the plate, $t = 6.0\mu\text{m}$ its thickness, $d = 0.50\text{mm}$ the slit separation, $D = 1.0\text{m}$ the slit-to-screen distance, and $\lambda = 500\text{nm}$ the wavelength of the light.

(a) The fringe-shift formula gives:

$$\Delta x = \frac{D\,(\mu - 1)\,t}{d} = \frac{(1.0\text{m})(0.50)(6.0 \times 10^{-6}\text{m})}{0.50 \times 10^{-3}\text{m}} = 6.0 \times 10^{-3}\text{m} = 6.0\text{mm}$$

The central bright fringe moves 6.0mm to the side on which the plate was placed.

(b) The fringe width of the apparatus, before any shift, is:

$$\beta = \frac{\lambda D}{d} = \frac{(500 \times 10^{-9}\text{m})(1.0\text{m})}{0.50 \times 10^{-3}\text{m}} = 1.0\text{mm}$$

The number of fringes that pass any fixed point on the screen as the plate is inserted is therefore:

$$N = \frac{\Delta x}{\beta} = \frac{(\mu - 1)\,t}{\lambda} = \frac{(0.50)(6.0 \times 10^{-6}\text{m})}{500 \times 10^{-9}\text{m}} = 6$$

Six bright fringes cross any fixed point on the screen during the insertion.

Making Sense of the Answer: *Notice that the answer to part (b) does not contain D or d at all; the number of fringes is set entirely by the extra optical path (μ − 1)t and the wavelength λ, both of which are properties of the light and the plate. This is one of the prettiest cancellations in the whole chapter. The geometric details of the apparatus, the very details examiners use to dress up the problem in different clothing, drop out the moment you ask for the number of fringes that have passed. The fringe count is, in this sense, more fundamental than the fringe shift: it depends only on physics, not on benchtop layout. The same plate inserted into a Newton's rings apparatus, or any other amplitude-division setup, would shift exactly the same number of fringes past any fixed observation point.*

Think Like a Physicist: *Inverting the relation N = (μ − 1)t/λ gives the experimental method that several Nobel-prize-quality measurements of refractive index and small thicknesses have rested on for over a century: count the fringes, know the wavelength, deduce the optical thickness. Whenever a question gives you fringe shift in fringe-counts rather than millimetres, this is the formula to reach for. Whenever you yourself want to measure the refractive index of an unknown thin film and you have a Young's apparatus on the bench, this is the experiment to perform.*

We have now completed the most important section of the chapter. Young's double-slit experiment, with its tightly packed bright and dark fringes a few millimetres apart on the wall, is the cleanest example of the split-source trick the previous section installed. The next section turns to its sibling, amplitude division, and treats the family of interference patterns produced by

thin films and air wedges. The geometry there is different, the formula 2μt for the optical path-difference replaces dy/D, but the physics is identical to what we have just developed: two coherent waves arrive at a single point, their phase-difference is set by a path-difference, and the bright-dark pattern follows by the same constructive-and-destructive logic that has now served us through two chapters of wave physics.

THIN FILMS AND AIR WEDGES

After Young's demonstration, with the sodium lamp still glowing and the blackboard still bearing the wave-counting arithmetic from the previous lesson, Mr. Akilikubwa walked to the laboratory tap, dipped a wire loop into the beaker of soap solution he had left out overnight, and lifted out a thin elastic disc of soap film stretched flat across the loop. He held it up under the desk lamp and waited.

The class watched. For two or three seconds, while the film was still everywhere thick, it appeared a vague pale grey. Then, as the soap drained downward under gravity, the film began to draw colour. Bands of colour started running across the disc, narrowest near the top, broadest near the bottom, slowly creeping downward as the film thinned: a violet just below the black band, then blue, then green, then yellow, then orange, then red, in unbroken sequence from the top of the rainbow to its bottom. At the very top of the loop, where the film had thinned the most, a band of pure black appeared, grew, and rose, until it dropped off the top edge and a fresh wave of colour took its place behind it.

Mr. Akilikubwa, still holding the loop: *"A soap film, lit by ordinary white light, behaves like a rainbow ironed flat and pinned to a wire frame: violet at the top of the colours, just below the black band; blue and green next; yellow, orange, and red at the bottom. Two questions, both of which this section will answer.* ***First, why colours at all?*** *Daylight is white. The film is colourless. The two together produce, in our window, every wavelength of the visible spectrum, neatly stratified by height like volcanoes by altitude. Where did the colours come from?* ***Second, why is the very top of the film black?*** *The black band sits exactly where the film is thinnest, just before it bursts. The thinner the film, the less light it should reflect at all;* ***why does the limit of vanishing thickness produce a perfectly dark reflection rather than a slightly faint one?"***

This section is about the answer. The answer is interference, but of the second kind installed in the coherence section: not the wavefront-division kind of Young's slits but the amplitude-division kind that splits a single wave into two by partial reflection at two parallel surfaces, lets the two parts travel along nearly identical paths through a thin layer, and recombines them at the eye of the observer. We have already seen the trick in outline. We now derive its formulas, apply them to soap films, to the air wedge between two clamped glass plates, and to the anti-reflective coatings that camera lenses depend on, and answer Mr. Akilikubwa's three questions exactly.

Amplitude division revisited

Picture the soap film of the demonstration, or any thin transparent film: a sheet of uniform thickness t and refractive index μ, with air of refractive index 1 above and below. Imagine a single ray of monochromatic light, of wavelength λ in vacuum, descending onto the top surface of the film from above. At that top surface two things happen at once, with the partial-reflection coefficient set by the refractive-index contrast across the boundary. Part of the wave reflects directly off the top surface and travels back up; call this reflection R_1. The remaining part of the wave refracts into the film, travels down through the film at the refraction angle, reaches the bottom surface, and there partially reflects again. The bottom-surface reflection travels back up through the film, refracts out at the top surface, and emerges into the air above moving parallel to R_1 but laterally displaced by a small amount; call this twice-bounced wave R_2.

The two waves R_1 and R_2 are coherent with each other (because they came from the same incident wave by amplitude division at the top surface) and they meet again at the eye or at the screen above the film. Whether the meeting is constructive (a bright fringe in the reflected light) or destructive (a dark fringe in the reflected light) is set, exactly as in Young's slits, by the phase difference between the two waves at the observation point. The phase difference here comes from two contributions: the geometric round-trip through the film, which depends on the film's thickness and its refractive index, and a phase shift at the top-surface reflection that wavefront division never had to worry about.

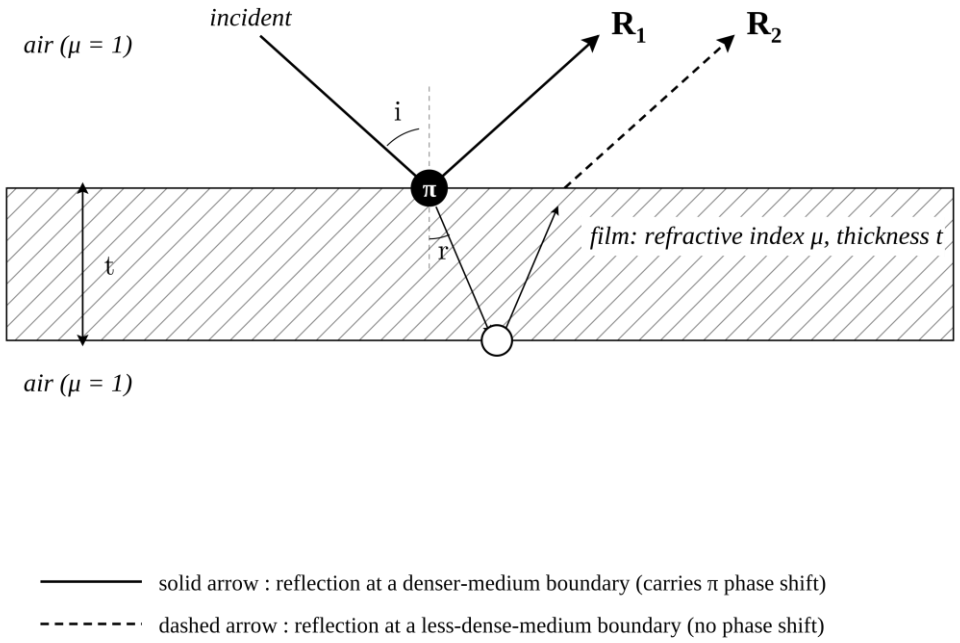


Figure: ***Amplitude division at a thin film.*** *The incident ray reflects partially at the top surface as R_1 (which suffers a π phase change because the film is denser than the air above it) and partially refracts into the film at angle r. The refracted ray reflects off the bottom surface without a phase change (the air beneath is less dense), returns to the top surface, refracts back out, and emerges as R_2 parallel to R_1 but laterally displaced.*

The phase change of π at the denser surface

One physical fact is owed before any formula. When a wave reflects off the boundary between two media, the reflected wave sometimes flips sign and sometimes does not, depending on which medium is denser. The rule, stated cleanly so it can be memorised in one sitting:

A wave reflecting off the surface of a medium that is optically denser than the one it came from suffers a phase change of π radians; equivalently, the reflected wave picks up an extra optical path of $\frac{\lambda}{2}$ relative to the incident wave. A wave reflecting off the surface of a medium that is optically less dense than the one it came from suffers no phase change at all; the reflected wave returns in phase with the incident wave.

The rule is not new. The same rule governed the reflection of a transverse wave on a stretched string from the previous chapter, where the wave was inverted at a tied-down end (a hard end) and was not inverted at a free end (a soft end). Light treats the denser medium as the hard end and the less-dense medium as the soft end, and the analogy is exact. Maxwell's equations supply the proof from the boundary conditions on the electric and magnetic fields, but the qualitative analogy with the string is enough for everything in this chapter.

Apply the rule now to our thin film. The wave R_1 reflects off the top surface of the film: it goes from air (less dense, $\mu = 1$) to film (denser, $\mu > 1$), and therefore picks up a phase shift of π, equivalently an extra $\frac{\lambda}{2}$ of optical path. The wave R_2 reflects off the bottom surface of the film: it is now inside the film going down, and reflects against the air beneath; light is going from the denser medium to the less dense, so no phase shift. Exactly one of the two reflections is hard-ended, and exactly one of the two waves carries an unavoidable extra $\frac{\lambda}{2}$ with it back to the eye. The asymmetry between the two reflections is what makes thin-film interference different from Young's interference, and it is responsible for every "surprise" the rest of this section will produce, starting with the dark cap on Mr. Akilikubwa's soap film.

The optical path difference through the film

Now the geometry. At near-normal incidence (which is the case we shall treat throughout, because it covers the soap film, the air wedge, and the anti-reflection coating all at once) the wave inside the film travels down through the film and back up, a geometric round-trip of total length 2t. But a metre of path inside a denser medium is not equivalent to a metre of path in air; the wavelength inside the film is shorter than the wavelength in vacuum by a factor of μ, and the optical path (measured in vacuum wavelengths) is therefore μ times the geometric path. The optical round-trip through the film is thus:

$$\text{Optical round} - \text{trip} = 2\mu t$$

(For a general angle of refraction r inside the film, the round-trip becomes $\mathbf{2\mu t \cos r}$; at near-normal incidence $\cos r \approx 1$ and the expression collapses to $2\mu t$.)

The total optical path difference between R_1 and R_2 at the observer is the optical round-trip through the film, plus the half-wavelength length-equivalent of the π phase shift at the top surface. We give the optical path difference the symbol Δx. In length units it reads:

$$\boldsymbol{\Delta x = 2\mu t + \frac{\lambda}{2}}$$

A word on terminology, *because two closely related quantities live in every interference calculation and they are not interchangeable. The optical path difference Δx just written down has units of length (metres). The phase difference Δφ at the observer has units of radians. The two are related by the master conversion from Chapter 7,* $\boldsymbol{\Delta\varphi = \frac{2\pi}{\lambda}\Delta x}$*, so that any constructive-or-destructive condition can be written in either form.* Throughout this section we work in length units, because the bright-fringe condition $\Delta x = n\lambda$ reads more cleanly than its phase-difference cousin $\Delta\varphi = 2n\pi$, but the two are saying the same thing.

The extra $\frac{\lambda}{2}$ in Δx is the bookkeeping signature of the asymmetric reflection. Without it, the bright and dark conditions on a thin film would be the constructive and destructive cases of a simple geometric path difference 2μt, with no surprises. With it, the conditions are systematically interchanged: what would have been the bright case is now the dark case, and vice versa. We turn the algebra next.

Conditions for bright and dark fringes

Constructive interference (bright fringe in reflected light) requires the total optical path difference to be an integer number of wavelengths:

$$\Delta x = 2\mu t + \frac{\lambda}{2} = n\lambda$$

Solving for the geometric part of the optical path:

$$\boldsymbol{2\mu t = \left(n - \frac{1}{2}\right)\lambda, \quad n = 1, 2, 3, \ldots}$$

Destructive interference (dark fringe in reflected light) requires the total optical path difference to be a half-integer number of wavelengths. Repeating the algebra with a half-integer on the right-hand side:

$$2\mu t + \frac{\lambda}{2} = \left(n + \frac{1}{2}\right)\lambda$$

$$\boldsymbol{2\mu t = n\lambda, \quad n = 0, 1, 2, \ldots}$$

Notice the value of n that the dark condition allows: $n = 0$, which corresponds to a film of vanishing thickness $t \to 0$. A film thinned all the way down to nothing exactly satisfies the dark-fringe condition. It reflects no light at any wavelength. The two reflections R_1 and R_2 from a film of zero thickness arrive at the eye almost coincident in space (because the bottom and the top surfaces have merged), almost coincident in time (because the round-trip through nothing takes no time), and exactly out of phase (because the top-surface reflection has flipped sign while the bottom-surface reflection has not). They cancel each other exactly, at every wavelength. That is the black band at the top of Mr. Akilikubwa's loop. The film has drained itself locally to a thickness much smaller than any visible wavelength, and the interference is destructive for every colour the eye can see. The film has become its own perfect anti-reflection coating, which is also the name we are going to give to the deliberate engineering version of this same physics three subsections from now.

Why the colour bands stack in the order they do

Below the black cap, the film has not yet drained to nothing. At each height up the loop, the film has some local thickness t set by gravity, surface tension, and the time the bubble has been hanging there. The bright-fringe condition can be solved at each height for the wavelength that is reinforced in the reflected light:

$$\lambda = \frac{2\mu t}{\left(n - \frac{1}{2}\right)}$$

With $n = 1$, the lowest non-trivial order, this gives $\lambda = 4\mu t$. For a soap film of refractive index roughly 1.33, the visible wavelengths between 400nm and 700nm are reinforced in the reflected light at thicknesses between 75nm and 132nm. As the film drains, the local thickness at any given height passes downward through this range, and the colour reinforced moves through the spectrum from red (700nm at 132nm thickness) to violet (400nm at 75nm thickness). The colours we see on the bubble are not the source colours but the wavelengths reinforced by the film thickness at each height, and they stack in the order set by the formula above.

The order is the reverse of the prism order, by the way: *a prism orders visible colours by their refractive indices, with red bent the least and violet bent the most. A soap film orders them by film thickness, with red at the thickest film and violet at the thinnest.* The fact that both orderings can be derived from the wave-mechanical account of light, by two completely different physical mechanisms (refraction in the prism, two-surface interference in the soap film), is one of the strongest internal-consistency arguments for the wave theory of light. Two roads, one map.

The Air Wedge

A particularly clean and quantitative version of the same physics is produced not by a soap film, which drains and bursts before any careful measurement can be performed, but by clamping two flat glass plates along one edge and inserting a thin spacer (a wire, a human hair, or a stack of foil) along the opposite edge. The two plates meet at the clamped edge with zero gap between them, and separate gradually as the spacer is approached, with the local thickness of the gap growing linearly along the wedge. The space between the two plates is air, and we therefore call the arrangement an air wedge.

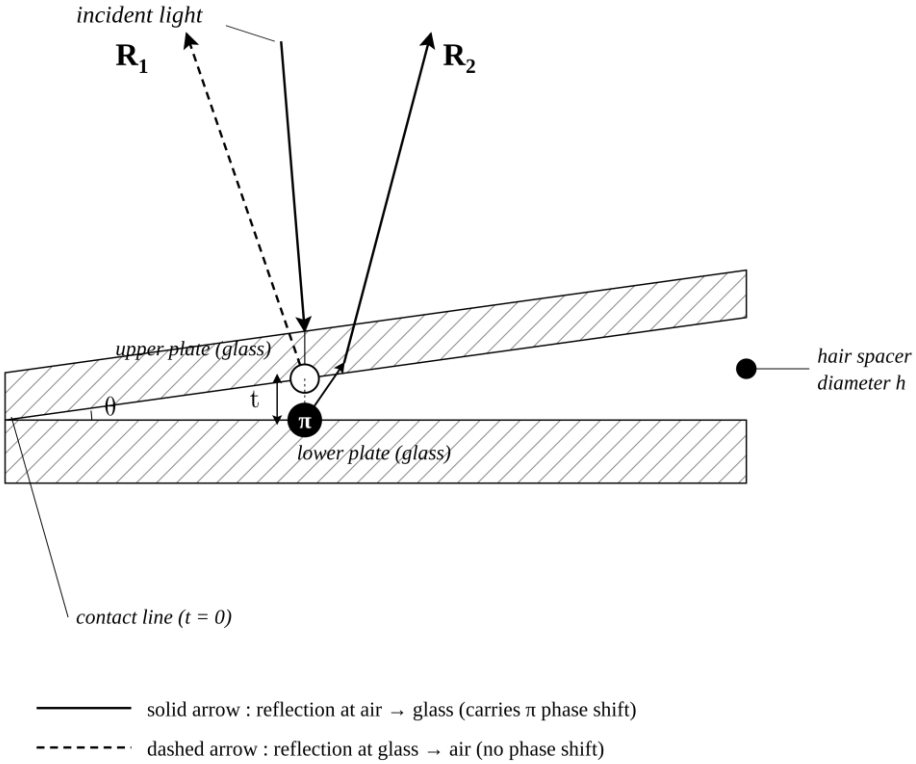


Figure: *An air wedge formed by two glass plates clamped at their left edge (the contact line, where t = 0) and separated at their right edge by a hair of diameter h. The local thickness t of the air gap grows linearly along the wedge. Light incident from above reflects off the bottom of the upper plate (no phase change, glass → air) and off the top of the lower plate (π phase change, air → glass). The two reflections interfere; the fringe spacing along the wedge encodes the wedge angle θ and the wavelength λ.*

The geometry is identical to the soap film with two small differences. **First,** *the refractive index of the air between the plates is* $\mu = 1$*, so the optical path through the wedge reduces from* $2\mu t$ *to simply* $2t$. **Second,** *the identification of which reflection picks up the π phase shift is reversed.* R_1, the wave that reflects from the lower surface of the upper glass plate, is going from glass (denser) to air (less dense): no phase shift, soft-end reflection. R_2, the wave that reflects from the upper surface of the lower glass plate, is going from air (less dense) to glass (denser): π phase shift, hard-end reflection. Exactly one of the two reflections is hard-ended, and the asymmetric phase contribution is the same $\frac{\lambda}{2}$ as before. The bright and dark conditions therefore look the same on the page as the soap-film conditions, with the refractive index dropped out:

$$\text{Bright (reflected): } \mathbf{2t = \left(n - \frac{1}{2}\right)\lambda, \quad n = 1, 2, 3, \ldots}$$

$$\text{Dark (reflected): } \mathbf{2t = n\lambda, \quad n = 0, 1, 2, \ldots}$$

At the clamped edge, where $t = 0$, the dark condition is satisfied with $n = 0$: the contact line appears as a dark fringe. This is the air-wedge analogue of the soap-film black cap, and it occurs for the same physical reason. As the thickness grows linearly along the wedge, the dark fringes occur at thicknesses $0, \frac{\lambda}{2}, \lambda, \frac{3\lambda}{2}, \ldots$, and the bright fringes interleave between them at the half-wavelength thicknesses.

Consecutive dark fringes are therefore separated by a thickness change of:

$$\Delta t = \frac{\lambda}{2}$$

The wedge angle θ (the small angle between the two plates, always small enough that $\tan\theta \approx \theta$ in radians) translates this thickness change into a horizontal distance β between consecutive fringes on the plate:

$$\theta\beta = \frac{\lambda}{2}$$

$$\theta = \frac{\lambda}{2\beta}$$

This is the air-wedge angle formula. It is the easiest experimental method for measuring a very small angle (the angle made by a slightly bent plate, or by a hair trapped between two slides) with a precision the naked eye could not approach.

A second corollary concerns the total number of fringes visible across the entire wedge. If the spacer at the open end has thickness h, the optical path at the open end is 2h, and the number N of dark fringes spanning the wedge from contact line to spacer is:

$$N = \frac{2h}{\lambda}$$

This formula provides one of the cleanest measurements of the wavelength of light a school laboratory can perform without buying any specialist instrument. A human hair, two glass slides, and a benchtop sodium lamp suffice. Count the dark fringes, measure the hair diameter with a micrometer screw gauge, divide twice the diameter by the count, and you have the wavelength of sodium light to the precision the integers allow.

Application: Anti-reflection coatings

A camera lens or a spectacle lens loses about four per cent of the light incident on it to reflection at the front surface. For a sophisticated zoom lens with twelve or fifteen elements, the cumulative loss adds up to a third or more of the incoming light, and the reflected fraction bounces around between elements until it eventually reaches the camera's sensor as a faint ghost image superimposed on the wanted one. Both effects are undesirable, and a single application of thin-film interference fixes both: deposit on the front surface of the lens a thin transparent coating of carefully chosen thickness, such that the wave reflected off the air-coating boundary interferes destructively with the wave reflected off the coating-glass boundary. The two reflected waves cancel each other; the lens reflects essentially no light at the design wavelength; the lost light is entirely transmitted into the glass as intended.

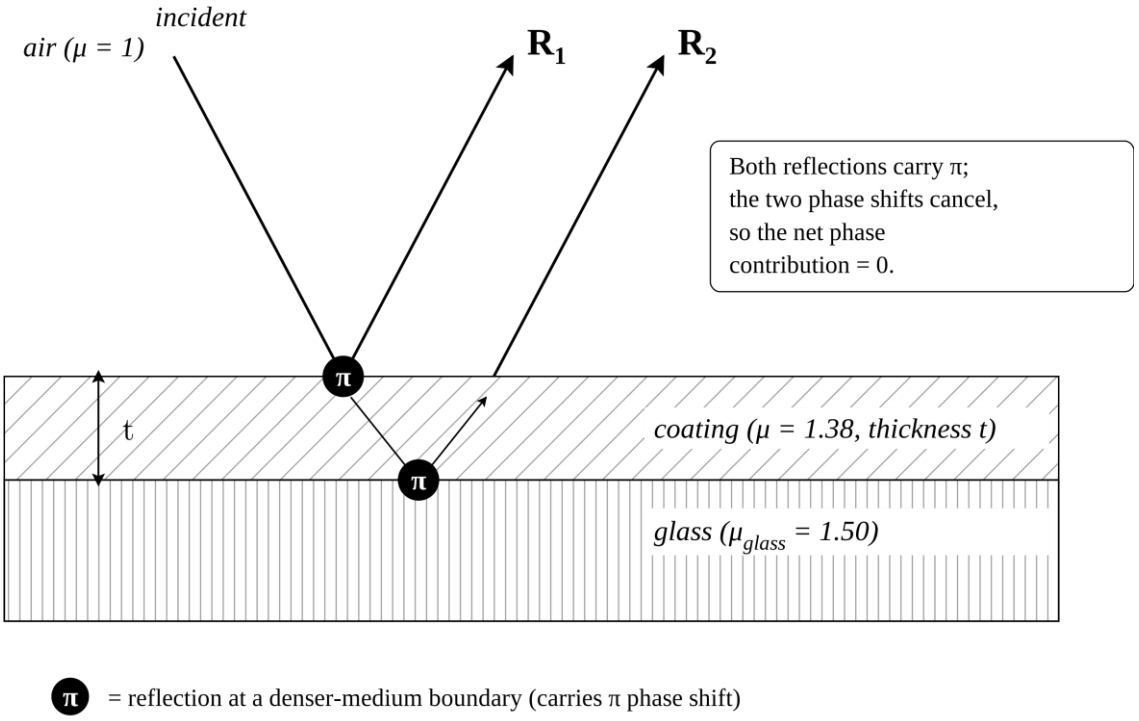


Figure: ***An anti-reflection coating of refractive index μ and thickness t deposited on a glass surface.*** *Both reflections (R_1 at the air-coating interface and R_2 at the coating-glass interface) suffer a π phase change because the medium beyond each interface is denser. The two phase changes cancel from the optical-path-difference bookkeeping, leaving the geometric round trip 2μt as the only contributor.*

The geometry is a thin film of refractive index μ and thickness t deposited on the front face of glass of refractive index μ_{glass}, with air of refractive index 1 on the outside. The two reflected waves are R_1 (off the air-coating boundary, where coating is denser than air, so a π phase change) and R_2 (off the coating-glass boundary, where glass is denser than coating, so another π phase change). Both reflections carry a π phase shift, and the two contributions cancel each other out of the difference. The condition for destructive interference in the reflected light is therefore the geometric condition alone:

$$2\mu t = \left(n + \frac{1}{2}\right)\lambda, \quad n = 0, 1, 2, \ldots$$

The thinnest coating that achieves this corresponds to n = 0:

$$2\mu t = \frac{\lambda}{2}$$

$$\mathbf{t = \frac{\lambda}{4\mu}}$$

This is the quarter-wave thickness for an anti-reflection coating, and it is the canonical formula for designing a single-layer coating optimised for one wavelength.

Two further refinements are worth knowing without proof, and we shall use one of them in the worked example below.

First, the cancellation between R_1 and R_2 is most complete when the coating has refractive index $\mu = \sqrt{\mu_{air}\mu_{glass}}$, because that condition makes the amplitudes of the two reflected waves equal as well as their phases opposite, so they cancel exactly rather than approximately. For air on crown glass at $\mu_{glass} = 1.50$, the ideal coating refractive index is $\sqrt{1.50} \approx 1.22$; the magnesium fluoride routinely used in commercial lenses has $\mu = 1.38$, which is close enough to give a residual reflected fraction of about 1.5 per cent compared with the uncoated 4 per cent, a substantial improvement but not a perfect cancellation.

Second, the formula $t = \frac{\lambda}{4\mu}$ optimises the coating for one wavelength only; wavelengths above and below it reflect slightly more than the uncoated lens, which is why coated lenses have a characteristic purple sheen in reflection (the green-yellow centre of the spectrum, at the design λ, is suppressed, leaving the blue and the red ends to dominate).

With the ideas now simmering nicely on the page, the next three worked examples are the spoon with which we taste them and feel the flavour of the physics for ourselves. Pen ready.

BINDER Example 10

Two flat glass plates are clamped along one edge to form an air wedge. At the open end, a fine human hair of diameter 0.060mm is inserted between them to act as a spacer. The wedge is illuminated normally with monochromatic light from a sodium lamp, and a travelling microscope counts a total of 240 dark fringes across the wedge from the contact end to the open end. Calculate the wavelength of the sodium light.

Solution

Let h denote the maximum thickness of the air wedge (the diameter of the hair) and N the total number of dark fringes counted from the contact end. The thickness step between consecutive dark fringes is $\frac{\lambda}{2}$, so the relation between h, N, and λ is:

$$N = \frac{2h}{\lambda}$$

Solving for the wavelength:

$$\lambda = \frac{2h}{N}$$

Substituting $h = 0.060\text{mm} = 6.0 \times 10^{-5}\ \text{m}$ and $N = 240$:

$$\lambda = \frac{2 \times 6.0 \times 10^{-5}\ \text{m}}{240} = 5.0 \times 10^{-7}\ \text{m} = 500\text{nm}$$

The wavelength of the sodium light is 500nm.

Making Sense of the Answer: *The answer lies in the green-yellow region of the visible spectrum, which is sensible for a school sodium lamp; the real sodium yellow doublet sits at 589nm in undisturbed atomic vapour, and our 500nm answer is in the right neighbourhood for a school sodium lamp whose lines have been broadened by a hot filament or by pressure-induced shifts. The arithmetic ran on three given numbers (the hair diameter, the fringe count, and the factor of two from the round-trip through the air gap) and produced an answer with five-percent precision without our owning a single specialist instrument.*

Think Like a Physicist: *Air-wedge fringe counts are the cleanest school measurement of the wavelength of light, because the only ingredients are a length (the diameter of the hair, measurable with a micrometer screw gauge to micrometre precision) and a count of integers (the number of fringes, requiring nothing more than patience and a travelling microscope). Neither quantity needs an oscilloscope, a photodetector, or a spectrometer. Notice that the formula has no spare parameters: everything that could carry a measurement uncertainty has been removed by the simple act of using a hair instead of a calibrated reference spacer. This is the cleanest kind of measurement physics knows how to do.*

BINDER Example 11

An air wedge is illuminated normally by monochromatic light of wavelength 600nm. A travelling microscope focused on the top surface of the wedge observes 11 consecutive bright fringes occupying a horizontal distance of 9.0mm along the wedge. Calculate the angle of the wedge in radians.

Solution

Let $\lambda = 600\text{nm}$. Eleven consecutive bright fringes have ten intervals between them (fringes are fenceposts, intervals are panels, and the panel count is always one less than the post count), so the fringe spacing is:

$$\beta = \frac{9.0\text{mm}}{10} = 0.9\text{mm} = 9.0 \times 10^{-4}\ \text{m}$$

The air-wedge angle follows directly from the fringe-spacing formula:

$$\theta = \frac{\lambda}{2\beta}$$

Substituting:

$$\theta = \frac{600 \times 10^{-9}\ \text{m}}{2 \times 9.0 \times 10^{-4}\ \text{m}} = 3.33 \times 10^{-4}\ \text{rad}$$

The wedge angle is 3.33×10^{-4} rad, which is about 0.019°.

Making Sense of the Answer: *Three hundred and thirty millionths of a radian is, in any normal sense, a perfectly absurd quantity to be able to measure. With a protractor on a sheet of paper, an angle of one one-thousandth of a degree would not be detected with the most generous reading. The fringe pattern, however, presents itself to the microscope as a series of dark and bright lines just under a millimetre apart, which any Form Five student with a steady hand can resolve to within a few percent. The optics has effectively multiplied the small angle of the wedge into a comfortably visible fringe spacing, and the multiplication factor is the fringe spacing divided by the wavelength of light,* $\frac{\beta}{\lambda} \approx 1500$ *in our case. Submicroscopic angles are made macroscopic by the wavelength of light; that is the price of admission to optical metrology.*

Think Like a Physicist: *The formula* $\theta = \frac{\lambda}{2\beta}$ *needed only the wavelength and the fringe spacing; nothing about the length of the wedge or the diameter of the hair appears. Whenever an answer turns out to be independent of a quantity the problem could in principle have asked you to use, that quantity is irrelevant to the physics of the result. Always look for what is missing from the answer as well as what is present; the absences are sometimes more informative than the presences.*

REAL Example 12

A camera-lens manufacturer in Mwanza wishes to apply a non-reflective coating to the front surface of a glass lens, with the goal of minimising reflected light at the green-yellow centre of the visible spectrum (around 550nm). The coating material chosen is magnesium fluoride, which has refractive index 1.38 and forms films of excellent uniformity by vacuum deposition. The glass has refractive index 1.50. Calculate the thinnest coating thickness that achieves destructive interference in the reflected light at the design wavelength, and comment briefly on the choice of magnesium fluoride as the coating material.

Solution

Let $\mu = 1.38$ be the refractive index of the coating, $\mu_{\text{glass}} = 1.50$ that of the glass beneath, and $\lambda = 550\text{nm}$ the design wavelength. Both reflections (air → coating and coating → glass) pick up a phase shift of π; the two phase shifts cancel each other out of the difference, and the condition for destructive interference in the reflected light is the purely geometric condition:

$$2\mu t = \left(n + \frac{1}{2}\right)\lambda$$

The thinnest coating corresponds to n = 0:

$$t = \frac{\lambda}{4\mu}$$

Substituting:

$$t = \frac{550 \times 10^{-9}\ \text{m}}{4 \times 1.38} = 9.96 \times 10^{-8}\ \text{m} \approx 100\text{nm}$$

The required coating thickness is about 100nm, or one tenth of a micrometre, a film thinner than a single bacterium and about a fifth the thickness of a soap-bubble band. Modern vacuum-deposition techniques achieve this routinely with thickness uniformity better than one nanometre across the lens.

On the choice of material: the cancellation between R_1 and R_2 is total only when their amplitudes are equal as well as their phases opposite. The phases are made opposite by the quarter-wave thickness condition above. The amplitudes are made equal by the additional condition that the refractive index of the coating should be the geometric mean of the refractive indices on either side, namely $\sqrt{\mu_{\text{air}}\mu_{\text{glass}}} = \sqrt{1.00 \times 1.50} \approx 1.22$. Magnesium fluoride at $\mu = 1.38$ is close to this ideal but not exactly on it; the residual reflected fraction at the design wavelength is about 1.5 per cent compared with the 4 per cent reflected by the uncoated glass. A reduction by a factor of roughly three, applied to every air-glass interface in a multi-element lens, is what makes modern professional camera lenses optically transparent enough that they appear almost black in reflection.

Making Sense of the Answer: *Around 100nm is the canonical thickness of a single-layer anti-reflection coating in the visible. The reason the answer is so robustly set by* $\frac{\lambda}{4\mu}$ *is that the formula has no free parameters once the design wavelength and the coating material are chosen; the engineering is entirely in the manufacturing precision (depositing exactly the right thickness, uniformly, across the lens), not in the design arithmetic. The difference of a few nanometres between the ideal-index coating and the nearly-ideal magnesium fluoride is absorbed in the residual 1.5 per cent reflection: a small price for the convenience of having a chemically robust and inexpensive material.*

Think Like a Physicist: *Whenever a problem mentions an anti-reflection coating, a quarter-wave plate, or any single-layer thin-film design, the formula* $t = \frac{\lambda}{4\mu}$ *is doing the work. Identify which surface gives which phase shift, decide whether you want destructive or constructive interference in the reflected light, and the integer-condition formula collapses to a single thickness for* $n = 0$*. Multi-layer coatings stack several quarter-wave layers of carefully chosen refractive indices, each designed to suppress the residual reflection of the layer below, and the same physics scales up by repetition: every layer contributes its own* $\frac{\lambda}{4\mu}$ *thickness condition, and the design problem reduces to choosing the indices.*

Amplitude division has now joined wavefront division in our toolkit. The new physical idea, the π phase change at the denser surface, has been spent three times: once on a thin film of uniform thickness, once on an air wedge of linearly varying thickness, and once on a quarter-wave anti-reflection coating of carefully chosen thickness. The next section turns to the most heavily examined apparatus that exploits amplitude division: Isaac Newton's rings, formed between a long-focus plano-convex lens and a flat glass plate. The geometry is a special case of the air wedge in which the gap thickness varies quadratically rather than linearly with horizontal position, producing circular fringes rather than straight ones. Everything else (the phase change, the dark central spot, the brightness ratios, the experimental protocol for measuring wavelengths to four significant figures with two pieces of glass and a microscope) follows by the same logic we have just spent.

NEWTON'S RINGS

After the air wedge of the previous section, one more arrangement of two reflecting surfaces deserves a section of its own. It is older than the soap film by almost two centuries (Isaac Newton himself reported it in 1675, and the 1704 publication of *Opticks* gave it the name it carries today), and it is by far the most heavily examined of all the thin-film phenomena in this chapter. The fringes are circular. They appear in a beautiful concentric ring pattern with a dark spot at the centre. The radii of the rings turn out to depend on the wavelength of light in a way clean enough to measure that wavelength to three significant figures with apparatus that fits on a school desk.

Mr. Akilikubwa, on the Tuesday after the air-wedge demonstration, had wheeled out from the laboratory preparation room a brass mount that held a plano-convex lens in light contact with a flat glass plate. He set a small sodium lamp to one side, clipped a thin glass square at forty-five degrees above the lens to act as a beam splitter, and mounted a travelling microscope vertically above the whole thing. He dimmed the windows, switched on the lamp, and called the class to the eyepiece.

What the students saw, once their eyes had adjusted to the small yellow disc of light reflected up to them, was not a uniform field. They saw a central dark spot the size of a chickpea, surrounded by a series of bright and dark concentric rings that shrank in spacing as one looked outward. The rings were the same yellow as the lamp, alternating with darkness, getting more and more crowded together until they faded into the edge of the field of view.

Kipanga: *Sir, what is producing the rings?*

Mr. Akilikubwa: *An air film. A wedge of air, the same kind we measured in the previous demonstration. But this time the wedge is not between two flat plates clamped at one end. It is between a flat plate and a curved lens resting on top of it. The air gap is zero where the lens touches the plate, and grows steadily as you move outward from the contact point. The fringes are the contours of constant thickness.*

That one sentence is the section. We will draw the apparatus, derive the relationship between the radius of a ring and the thickness of the air film beneath it, identify the optical condition for each ring to be dark or bright, read the wavelength of the source from the rings, and end with three small modifications to the experiment that examiners delight in asking about. Have your pen ready.

The apparatus and what the eye sees

The arrangement is shown in the next figure. A plano-convex lens of large radius of curvature sits with its curved side down on a flat glass plate, often called an optical flat. A thin film of air is therefore trapped between the underside of the lens and the upper face of the plate. Directly above the lens, a half-silvered glass square is tilted at forty-five degrees so that monochromatic light arriving horizontally from the right is partly reflected downward onto the lens, and the reflected light coming back up from the air film is partly transmitted upward into a low-power microscope mounted vertically above. The microscope magnifies the ring pattern onto your eye.

Apparatus for the observation of Newton's rings

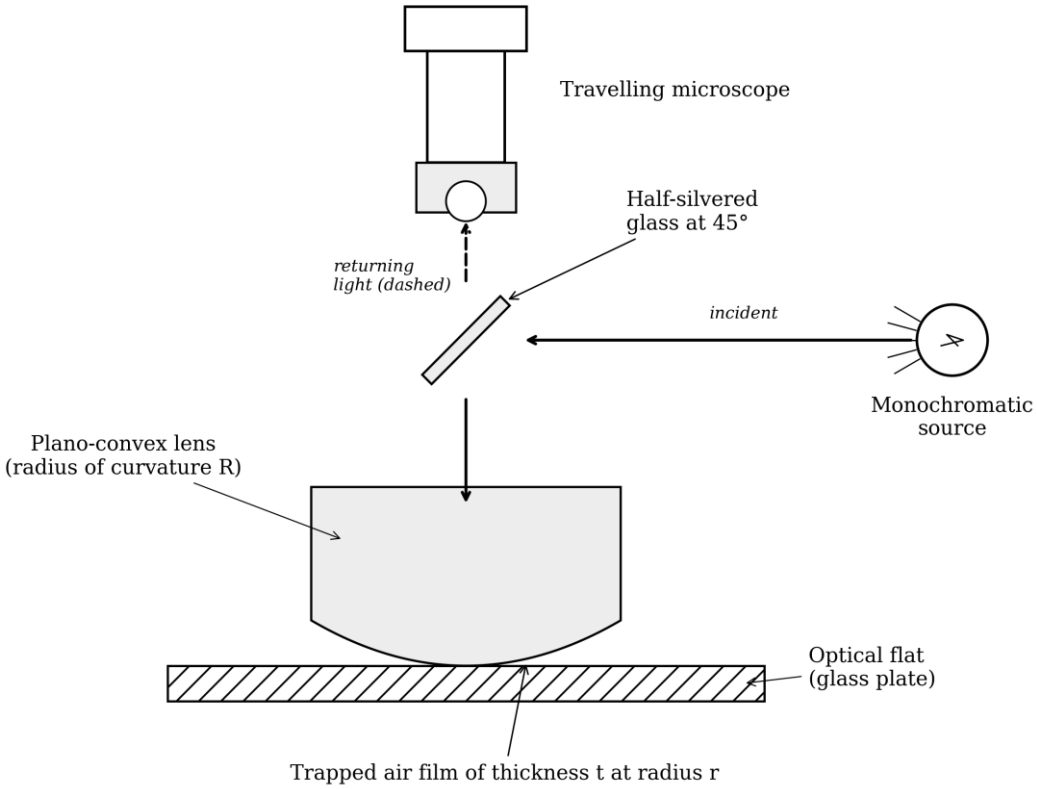


Figure: ***A plano-convex lens rests on an optical flat.*** *Light from a monochromatic source is bent downward onto the lens by a half-silvered glass at 45°. Two reflections (one from the lower curved surface of the lens, one from the upper surface of the plate) travel back up through the beam splitter and into the travelling microscope. The interference between these two reflections is what produces the ring pattern.*

The light striking the lens splits into two reflections of interest. The first, R_1, reflects off the lower curved surface of the lens (a glass-to-air boundary, with the light moving from glass into air). The second, R_2, reflects off the upper surface of the flat plate (an air-to-glass boundary, with the light moving from air into glass). The two reflected waves then travel back up to the eyepiece. Whatever pattern of constructive and destructive interference these two reflections produce, the microscope shows it to you.

The radius of curvature of the lower face of the lens, which we shall call R, is large (typically between half a metre and two metres in a school setup). The thickness of the air film at the very centre, where the lens touches the plate, is essentially zero. At a horizontal distance r from the contact point, the thickness t of the air film is determined by the geometry of the spherical lens surface. The next sub-section works out exactly what that thickness is.

A practical word before any algebra: The contact between the lens and the plate must be gentle and the surfaces must be clean. A grain of dust between them will tilt the lens slightly and skew the ring pattern into ovals. A fingerprint on either glass will paint a haze of additional fringes from the oil. The careful experimenter wipes both surfaces with lens tissue, places the lens on the plate without pressing, and accepts that the central dark spot will sometimes refuse to sit perfectly in the middle of the field of view. The geometry of the rings does not depend on the contact being in the centre of the microscope's field, so all is well.

Why the centre is dark

Before any algebra at all, one observation deserves attention. At the very centre, where the lens touches the plate, the air film has zero thickness. The two reflected waves R_1 and R_2 have travelled the same geometric distance. The naive prediction is that they should arrive in phase, interfere constructively, and produce a bright spot. The naive prediction is wrong. The centre is dark in every Newton's-rings experiment ever performed, and the reason is the same reason a soap film looks black at the top: an asymmetric phase change at reflection.

Kipanga: *Sir, the lens and the plate are touching at the centre. The air gap is zero. Zero path difference. Surely that means constructive interference; the centre should be bright.*

Mr. Akilikubwa: *That is exactly the trap, Kipanga. The geometric path difference at the centre is zero. But you have counted only one of the two ingredients. The other ingredient is the phase change of π that happens at every reflection off a denser medium. Walk through it once carefully.*

Kipanga: *The first reflection is off the bottom of the lens. The wave is inside the glass and reflects off the air gap. Going from denser to less dense. No phase change.*

Mr. Akilikubwa: *Correct.*

Kipanga: *The second reflection is off the top of the glass plate. The wave is inside the air and reflects off the glass plate below. Going from less dense to denser. Phase change of π.*

Mr. Akilikubwa: *Correct. So at the centre, your two reflections have a geometric path difference of zero and a phase difference of π. Reading this as an optical path difference, you have an extra λ/2 between them. They arrive perfectly out of phase. Destructive interference. The centre is dark.*

Kipanga: *All because one of the two reflections happens at a denser interface and the other does not.*

Mr. Akilikubwa: *All because of that. The asymmetry of the phase change is the whole story. Anyone who remembers only the geometric path difference will predict a bright central spot, and will then have to explain why their prediction disagrees with what the microscope clearly shows. The phase change is not optional.*

The argument generalises. At any point of the air film, the optical path difference between R_1 and R_2 is the geometric round-trip through the air (which is $2t$, because the wave goes down through the air and back up), plus a half-wavelength contribution from the single π phase change at the lower surface. In symbols,

$$\Delta x = 2t + \frac{\lambda}{2}$$

Dark fringes occur where the total optical path difference is a half-integer number of wavelengths (so the two waves arrive π out of phase), and bright fringes occur where it is an integer number of wavelengths (so they arrive in phase). Inserting the half-wavelength contribution and rearranging,

$$\text{dark: } 2t = n\lambda, \quad n = 0, 1, 2, 3, \ldots$$

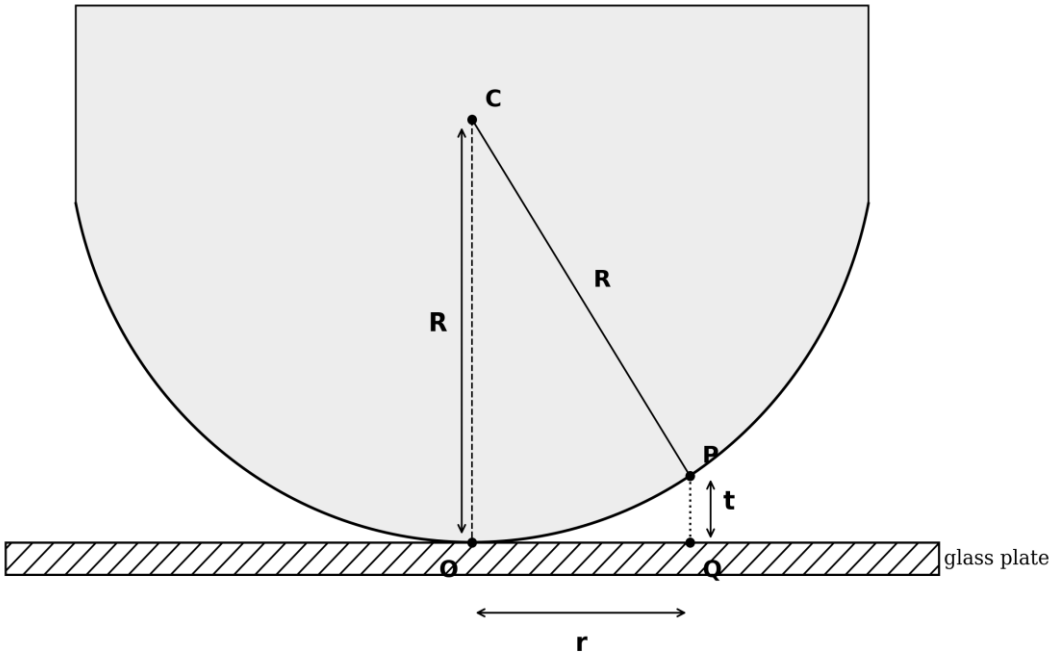
$$\text{bright: } 2t = (n - ½)\lambda, \quad n = 1, 2, 3, \ldots$$

The $n = 0$ case of the dark condition is the centre, where $t = 0$ satisfies $2t = 0 \times \lambda$ exactly. The first non-trivial dark ring occurs when $2t = \lambda$, the second when $2t = 2\lambda$, and so on. Each successive dark ring sits on a higher contour of the trapped air film.

The geometry that links radius to thickness

To turn the optical condition into a numerical formula for the radius of a ring, we need a relation between the thickness t of the air film and the horizontal distance r from the contact point. The relation is set entirely by the shape of the lens, and it is one of the cleanest pieces of geometry in the whole chapter.

Geometry of the trapped air film



From the right triangle: $r^2 + (R - t)^2 = R^2$, so $r^2 \approx 2Rt$ when t is much less than R

Figure: ***Cross-section of the lens on the plate.*** *The lower face of the lens is part of a sphere of radius R centred at C, sitting directly above the point of contact O. A general point P on the lens at horizontal distance r from the axis has the air film of thickness t beneath it (down to Q on the plate). The right triangle PQC immediately yields $r^2 + (R - t)^2 = R^2$, and dropping the t^2 term (because t is much smaller than R) gives $r^2 \approx 2Rt$.*

Place the centre of the spherical lens surface at the point C, directly above the contact point O by a distance equal to the radius of curvature R. Take any general point P on the lower face of the lens, at horizontal distance r from the axis. Drop a perpendicular from P down to the plate, meeting it at Q. The length PQ is the thickness t of the air film at that radial position.

The triangle PQC is right-angled at Q. Its horizontal side is r. Its vertical side is $R - t$ (the distance from Q up to C is the full radius minus the thickness of the gap). Its hypotenuse is CP, which equals R because P lies on the sphere. Pythagoras gives:

$$r^2 + (R - t)^2 = R^2$$

Expand the bracket and cancel the common R^2 from both sides:

$$r^2 + R^2 - 2Rt + t^2 = R^2,$$
$$r^2 = 2Rt - t^2$$

The term t^2 is utterly negligible. The thickness t of the air film is at most a few micrometres in a typical Newton's-rings setup, while the radius of curvature R is of order one metre. The ratio t/R is therefore about one in a million; the ratio t^2/R is one in a million times one micrometre, which is utterly invisible. We drop t^2 without losing anything that mattered, and we are left with the master geometric result of this section:

$$\mathbf{r^2 = 2Rt} \quad \text{equivalently} \quad \mathbf{2t = \frac{r^2}{R}}$$

Two sentences tell the whole story.

1) *The thickness of the air film, at horizontal distance r from the contact point, is proportional to the square of r.* The constant of proportionality is $1/(2R)$, which depends only on how curved the lens is.
2) *A flatter lens (larger R) gives a slower-growing air film and therefore wider-spaced rings.* A more strongly curved lens (smaller R) gives a faster-growing air film and tightly-packed rings. The rings widen and narrow with the geometry of the lens you happen to have chosen, but their structure follows a single universal formula.

The radii of the dark and bright rings

We now combine the optical condition for a dark fringe ($2t = n\lambda$) with the geometric relation ($2t = r^2/R$). Equating the right-hand sides gives the radius of the n-th dark ring:

$$\mathbf{r_n^2 = n\lambda R}, \quad \mathbf{n = 0, 1, 2, 3, \ldots}$$

Working in diameters rather than radii is often more convenient, because the travelling microscope reads the position of a ring on both sides of centre and the difference is the diameter directly. Doubling r gives $D_n = 2r_n$, and squaring yields:

$$\mathbf{D_n^2 = 4n\lambda R}$$

For bright rings, the optical condition is $2t = (n - ½)\lambda$. Combining with the same geometric relation gives the bright-ring radius and diameter as:

$$\mathbf{r_n^2 = (n - ½)\lambda R, \quad D_n^2 = (2n - 1)2\lambda R, \quad n = 1, 2, 3, \ldots}$$

Three features of these formulas earn a moment of attention before we put them to work.

First, *the diameters do not grow linearly with the ring number; they grow with the square root of n*. The first ring is one unit out; the fourth ring is two units out; the ninth ring is three units out. The rings crowd together as you move outward, which is what the eye saw at the microscope.

Second, *the radii depend on three quantities only: the integer n, the wavelength λ, and the radius of curvature R*. Nothing else in the geometry of the apparatus enters.

Third, *by measuring one ring and knowing two of the three quantities, you can compute the third.* That is the entire experimental method of Newton's rings: pick the unknown you want, build an apparatus that supplies the other two, and read the answer off the microscope.

Reading the wavelength from a graph

Measuring a single ring is risky in practice. The very innermost rings are crowded, broad, and hard to centre under the microscope's cross-hairs. The very outermost rings are dim because the intensity of the source falls off and the contrast washes. The reliable laboratory method is therefore to measure several rings, plot their squared diameters against the ring number, and extract the wavelength from the slope of the resulting straight line.

From the dark-ring formula $D_n^2 = 4n\lambda R$, a plot of D_n^2 against n is a straight line through the origin with slope $4\lambda R$. The wavelength is therefore:

$$\lambda = \frac{\text{slope}}{4R}$$

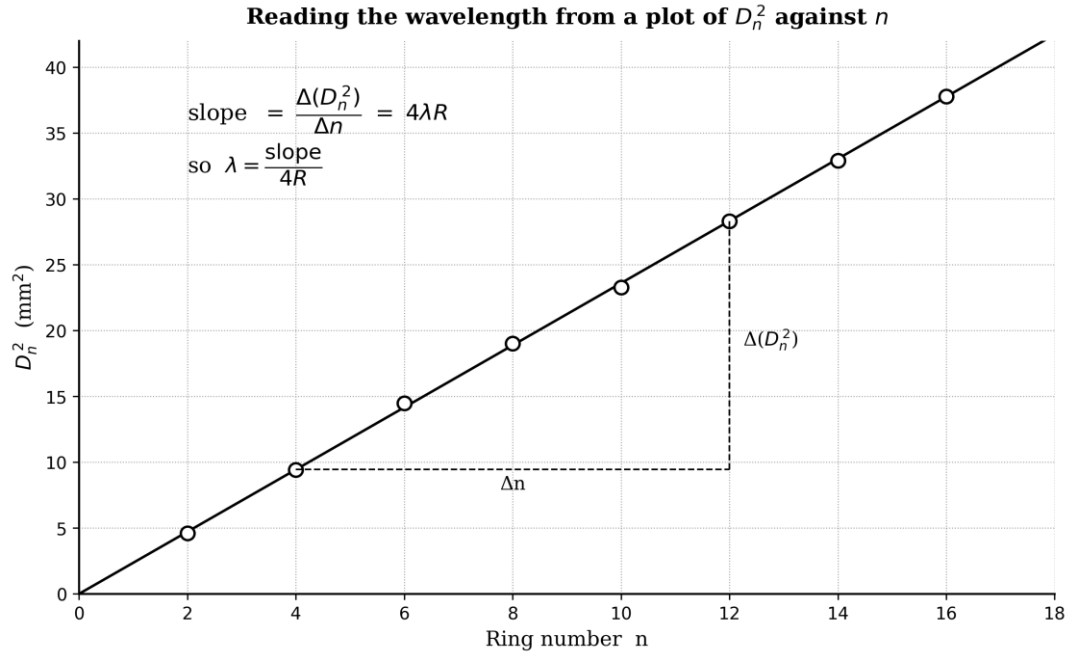


Figure: ***Squared diameter D_n^2 of the n-th dark ring plotted against the ring number n.*** *Each open circle is one measurement; the solid line is the best fit through the origin. The slope of the line, read off as $\Delta(D_n^2)/\Delta n$ from the dashed right triangle, equals $4\lambda R$, so the wavelength of the source is λ = slope / (4R).*

Two advantages of the graphical method are worth pinning down.

First, *the slope is the average rate of change over many rings, so any small error in counting which ring is which (off by one at the centre, where the rings are crowded and hard to count) averages out.*

Second, *the slope does not depend on whether the n = 0 ring is at the centre of the microscope's cross-hairs or slightly displaced (the slope is invariant under a constant offset in n).*

The graphical method is therefore robust against the very mistakes a careful experimenter is most likely to make. Single-ring measurements are usable but always less trustworthy than a slope reading from five or more rings.

Filling the gap with a liquid

So far the gap between the lens and the plate has been air, with refractive index $\mu = 1$. A favourite extension of the experiment, both in textbooks and in physics laboratories interested in measuring the refractive index of clear liquids, is to pour a small drop of liquid between the lens and the plate and re-record the rings.

With a liquid of refractive index μ filling the gap, the optical round-trip through the gap is no longer $2t$ but $2\mu t$ (each metre of liquid contributes μ metres of optical path, because the wavelength of light shortens to λ/μ inside the liquid). The π phase change at the air-to-glass or liquid-to-glass lower boundary is still there provided the liquid is less dense than the glass plate (which is true for water, glycerine, alcohol, and almost every transparent liquid a school laboratory will ever contain). The optical condition for dark fringes becomes:

$$\mathbf{2\mu t = n\lambda},$$

and combining with the geometric $2t = \frac{r^2}{R}$ gives the new ring radius:

$$\mathbf{r_n^2 = \frac{n\lambda R}{\mu}}, \quad \mathbf{D_n^2 = \frac{4n\lambda R}{\mu}}$$

The rings shrink by a factor of the square root of μ everywhere. The ratio of the diameter of any given ring in air to the diameter of the same ring in liquid is therefore exactly the square root of μ, and squaring this ratio gives the refractive index directly:

$$\mathbf{\mu = \frac{D_n^2(air)}{D_n^2(liquid)}}, \quad \text{at the } \mathbf{same} \text{ ring number n.}$$

Two points of practical good sense.

1) The same formula holds if you compare the slopes of the two D_n^2 against n plots, one for air and one for liquid:

$$\mu = \frac{\text{slope(air)}}{\text{slope(liquid)}},$$

which is the graphical version and is more accurate.

2) The result is independent of which particular ring you happen to compare, which is a useful self-check on the experiment: if your computed μ comes out different for the fourth ring than for the fourteenth ring, something has gone wrong (probably air has been re-admitted into the gap, or the liquid has not fully replaced the air).

Three small changes that the rings respond to

Three simple modifications of the apparatus produce qualitatively different ring patterns, all of which can be predicted from the physics already in your head. Examiners are fond of these because they reward the student who actually understands the concept and punish the student who has memorised $D_n^2 = 4n\lambda R$ without thinking about where it comes from.

Modification 1: Replace the sodium lamp with a white-light source

Each wavelength now produces its own family of rings, with red rings of any given order at slightly larger radii than violet rings of the same order (because the diameter is proportional to the square root of the wavelength). The central spot remains dark for every wavelength simultaneously, because the half-wavelength contribution from the single π phase change cancels every wavelength's geometric zero at $t = 0$. A few orders out from the centre, the wavelengths still produce distinguishable coloured rings (violet on the inside, red on the outside, for each order). Further out, the rings of different colours start to overlap, the colours mix, and after perhaps half a dozen orders the pattern washes into a uniform pale grey. The following figure shows both pictures side by side.

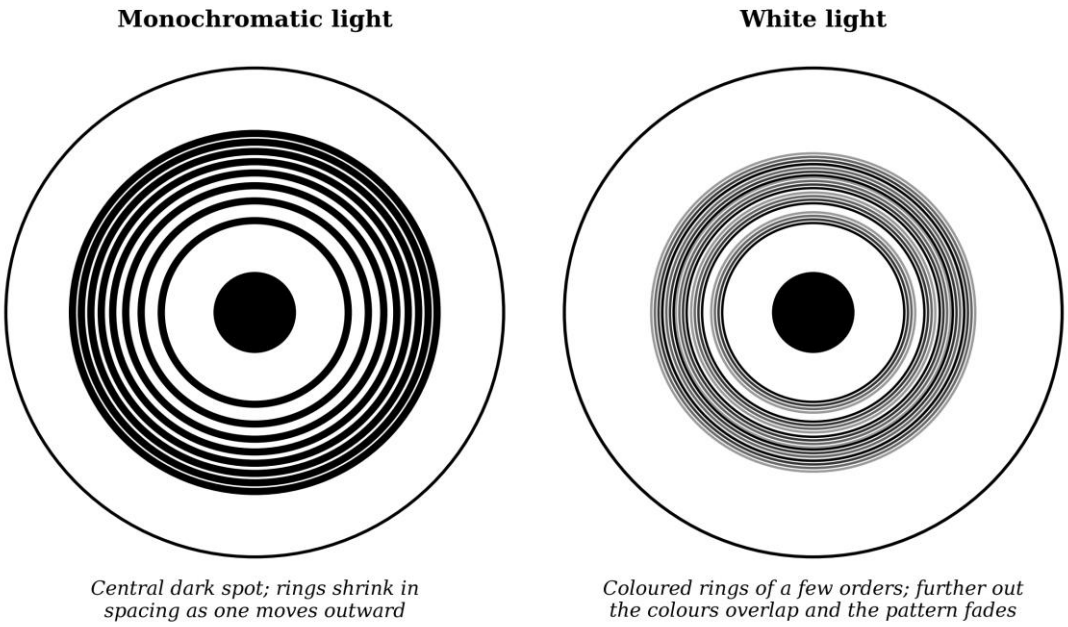


Figure: *Left: the ring pattern observed under sodium light, with a central dark spot and outward rings that crowd together as 1/√n. Right: the same apparatus illuminated by white light, with coloured rings of a few orders before different wavelengths overlap and the pattern fades.*

Modification 2: Silver the upper surface of the flat plate, so that the lower reflection becomes a metallic reflection rather than a glass reflection

A metallic reflection also picks up an extra π phase change. Both reflections now suffer the same phase shift, the asymmetry that produced the dark central spot disappears, and the geometric condition for constructive and destructive interference reverses. At the centre, $t = 0$ with no net phase difference between the two waves now gives constructive interference: the central spot becomes **bright**. The pattern outside the centre is the same alternating sequence of rings, but interchanged: dark where it used to be bright, bright where it used to be dark.

Modification 3: Introduce a few drops of a transparent liquid between the lens and the plate

This is the case treated in the previous sub-section. Every ring shrinks in diameter by a factor of $\sqrt{\mu}$, the spacing between consecutive rings contracts by the same factor, and the central spot remains dark (the optical condition $2\mu t = 0$ plus the half-wavelength is still destructive). A side-by-side observation of the same ring before and after the liquid is added gives the refractive index of the liquid as the ratio of squared diameters.

A small warning is necessary for modification 2: The phase change on metallic reflection is not exactly π for all metals and all wavelengths. Silver at sodium yellow comes very close to π, which is why the silvered-plate experiment is described as straightforwardly reversing the pattern. For other metals and other colours, the central spot becomes bright but the contrast of the surrounding rings can be diminished from what the air-gap experiment delivers. Real laboratory work always uses silver.

With this solid theoretical foundation in place, we are now ready to see these ideas come to life through four worked examples.

BINDER Example 13

In a Newton's-rings setup the radius of curvature of the plano-convex lens is 0.80m. The radius of the eighth dark ring is measured under the travelling microscope and found to be 1.94mm. Calculate the wavelength of the light used.

Solution

Let the wavelength be λ, the ring number $n = 8$, the radius of curvature $R = 0.80\text{m}$, and the radius of the eighth dark ring $r_8 = 1.94 \times 10^{-3}\text{m}$. Apply the dark-ring formula:

$$r_n^2 = n\lambda R$$

Rearrange for λ:

$$\lambda = \frac{r_n^2}{nR}$$

Insert numbers with units:

$$\lambda = \frac{(1.94 \times 10^{-3}\text{m})^2}{(8 \times 0.80\text{m})} = 5.88 \times 10^{-7}\text{m}$$

The wavelength of the light used is $5.88 \times 10^{-7}\text{m}$ or 588nm.

Making Sense of the Answer: *588nm sits in the yellow part of the visible spectrum, just a hair below the famous sodium D-doublet at 589.0nm and 589.6nm. Since the apparatus used a sodium lamp, the answer practically tells us we used a sodium lamp, which is exactly the calibration we should hope for from a well-set-up experiment. If our answer had come out at, say, 280nm or 1300nm, we would know without further checking that something was wrong: 280nm is invisible ultraviolet and 1300nm is infrared, and the lamp on the bench was glowing visibly yellow. Use this kind of physical reality-check at the end of every calculation in this chapter; the visible spectrum is roughly* ***400nm to 700nm****, and any answer outside that range demands an explanation.*

Think Like a Physicist: *When you are given a single ring radius and asked for the wavelength, the only fact that matters is $r_n^2 = n\lambda R$. Plug, chug, done. The graphical method (slope of D^2 against n) is more accurate when you can measure several rings, but for a school problem where only one ring is given, the single-ring formula is the right tool. Pick the formula that matches the number of data points you have been given.*

REAL Example 14

In a Newton's-rings experiment with a lens of radius of curvature 1.00m, the diameters of the sixth and twelfth dark rings are measured to be 3.76mm and 5.32mm respectively. Calculate the wavelength of the light used.

Solution

Let D_6 and D_{12} be the two ring diameters, $R = 1.00\text{m}$ the radius of curvature, and λ the unknown wavelength. From $D_n^2 = 4n\lambda R$, the squared diameter grows linearly with n, so the difference of the squared diameters of two rings depends only on the difference of the ring numbers:

$$D_{12}^2 - D_6^2 = 4(12 - 6)\lambda R = 24\lambda R$$

Rearrange for λ:

$$\lambda = \frac{(D_{12}^2 - D_6^2)}{24R}$$

Insert numbers with units:

$$\lambda = \frac{(5.32 \times 10^{-3}\text{m})^2 - (3.76 \times 10^{-3}\text{m})^2}{24 \times 1.00\text{m}} = 5.90 \times 10^{-7}\text{m} = 590\text{nm}.$$

The wavelength of the light is approximately $5.90 \times 10^{-7}\text{m}$.

Making Sense of the Answer: *590nm is right in the middle of the yellow part of the spectrum and is a sodium lamp again. Notice that the difference-of-squares method has not asked us to know the absolute ring number of either ring; it has only asked us to know that the two rings are separated by six orders. If our count of which ring was the sixth happened to be off by one at the centre (the centre is crowded and easy to miscount), the answer would shift by exactly nothing, because the labels n and n + 6 would have shifted together. The method is far more robust than picking a single ring's radius.*

Think Like a Physicist: *Whenever a problem gives you two ring measurements, do not solve it as two single-ring problems. Subtract the squared diameters, divide by the difference of ring numbers times 4R, and read the wavelength off. The same arithmetic with three or more ring measurements becomes a slope on a plot, which is even more reliable. Always prefer differences of squared diameters to single radii.*

REAL Example 15

In the apparatus of the previous example, the air gap between the lens and the plate is now filled with a clear liquid. The twelfth dark ring, which had a diameter of 5.32mm in air, now has a diameter of 4.60mm with the liquid present. The radius of curvature of the lens is unchanged at 1.00m and the source is the same sodium lamp. Calculate the refractive index of the liquid.

Solution

Let $D_{12}(\text{air})$ and $D_{12}(\text{liq})$ be the two diameters of the same ring before and after the liquid is added, and μ the refractive index of the liquid. The ring formula in the liquid is $D_n^2 = 4n\lambda R/\mu$, and the air formula is $D_n^2 = 4n\lambda R$, so dividing the air diameter squared by the liquid diameter squared at the same ring number gives μ directly:

$$\mu = \frac{D_n^2(\text{air})}{D_n^2(\text{liq})}$$

Insert numbers with units:

$$\mu = \frac{(5.32 \times 10^{-3}\text{m})^2}{(4.60 \times 10^{-3}\text{m})^2} = 1.338$$

The refractive index of the liquid is approximately 1.34.

Making Sense of the Answer: *A refractive index of 1.34 is essentially the refractive index of fresh water (which is 1.33 at room temperature for yellow light). If the answer had come out at, say, 0.85 (smaller than 1), we would know without checking that the calculation had gone wrong: a refractive index less than one does occur for some materials at some wavelengths, but it is not on the menu for ordinary liquids in a school laboratory. If the answer had come out at 2.5, we would suspect a*

*transcription error: even diamond, the densest optical solid commonly handled, is at 2.42 in visible light. Trust the answer if it falls between **1.3 and 1.7**, sigh at the answer if it falls outside, and re-do the arithmetic.*

Think Like a Physicist: *Whenever you compare a ring diameter in two different media at the same ring number, square the ratio of the diameters and you have the refractive index. Whenever you compare ring diameters between two different ring numbers in the same medium, take the difference of the squares (Example 14). The two manoeuvres are the two faces of the same algebra, and a Newton's-rings problem is almost always one or the other.*

HOT Example 16

A Newton's-rings apparatus is illuminated by light of wavelength 600nm. When the gap between the plano-convex lens and the flat glass plate contains air, the diameters of the fourth and twelfth dark rings are measured to be 3.80mm and 6.58mm respectively. After the same gap is filled with a transparent liquid, the diameters of the same two rings become 3.10mm and 5.37mm respectively. Calculate (a) the radius of curvature of the lens and (b) the refractive index of the liquid.

Solution

Let R be the unknown radius of curvature of the lens and μ the unknown refractive index of the liquid. We have four pieces of data: two ring diameters in air and two ring diameters in liquid, all at $\lambda = 6.00 \times 10^{-7}\text{m}$. The two unknowns require two independent equations.

Build the first equation from the air measurements. The difference of squared diameters of the twelfth and fourth dark rings in air is $D_{12}^2(\text{air}) - D_4^2(\text{air}) = 4 \times (12 - 4) \times \lambda R = 32\,\lambda R$:

$$D_{12}^2(\text{air}) - D_4^2(\text{air}) = 32\,\lambda R$$

$$R = \frac{D_{12}^2(\text{air}) - D_4^2(\text{air})}{32\,\lambda} = \frac{(6.58 \times 10^{-3}\text{m})^2 - (3.80 \times 10^{-3}\text{m})^2}{32 \times 6.00 \times 10^{-7}\text{m}} = 1.50\text{m}$$

The radius of curvature of the lens is 1.50m.

Build the second equation from the liquid measurements. The same difference, in the liquid is:

$$D_{12}^2(\text{liq}) - D_4^2(\text{liq}) = \frac{32\,\lambda R}{\mu}$$

$$\mu = \frac{32\,\lambda R}{D_{12}^2(\text{liq}) - D_4^2(\text{liq})} = \frac{32 \times 6.00 \times 10^{-7}\text{m} \times 1.50\text{m}}{(5.37 \times 10^{-3}\text{m})^2 - (3.10 \times 10^{-3}\text{m})^2} = 1.50$$

The refractive index of the liquid is approximately 1.50, which is consistent with a clean cedar oil or a light mineral oil at room temperature.

Making Sense of the Answer: *Both answers carry weight checks. The radius of curvature of 1.50m is in the right range for a school lens (which is typically 0.5m to 2m); a lens with R = 1.50m has a very shallow curvature, which is exactly what you want for Newton's rings because shallow lenses give wide, well-spaced rings that the microscope can resolve. The refractive index of 1.50 is too high for water (which is 1.33) and too low for diamond (which is 2.42), so it is something organic and oily. Cedar oil at the sodium D-line happens to be 1.516, mineral oil sits between 1.46 and 1.50, and most clear cooking oils are in the same range.*

Think Like a Physicist: *Whenever a problem hands you two unknowns and four measurements, look for two equations. Here the two equations both came from the difference-of-squares manoeuvre, once in air and once in liquid; their ratio cleanly delivered μ without even substituting R, and the absolute value of either delivered R if you wanted it. The general rule is: more measurements than unknowns is a gift, not a complication. Always set up the simultaneous equations first, then solve them. Plugging numbers into a single formula before you have written the system on paper is a common way to drop into algebraic confusion.*

We have spent the air-film geometry three times: on a thin film of uniform thickness, on an air wedge of linearly varying thickness, and now on the thicker-at-the-edge spherical air film of Newton's rings. The same two ingredients, the geometric round-trip $2t$ through the gap and the asymmetric phase change of π at one of the two boundaries, have predicted the dark-fringe-then-coloured-fringes appearance of a soap film, the parallel fringes of an air wedge, and the concentric ring pattern of Newton's apparatus. The wavelength measurement methods are not three different methods either: they are three geometrical dressings of the same idea, namely that a known geometric distance times a known optical condition gives a known wavelength.

Two more interference systems would fit in this section, but the chapter has work to do elsewhere and we shall not stay any longer. The **Michelson interferometer** and the **Fabry-Perot etalon** both repeat the same physics with different geometry; the first is the subject on the speed of light, and the second is the workhorse of laser-quality optical filters. They are flagged here so that the reader who meets them later recognises the family resemblance. For now we close interference and move to the next physical question of the chapter: *when a single wave of light passes a single obstacle whose size is comparable with the wavelength, what does the wave do?*

DIFFRACTION

The first half of the chapter belonged to interference. Two waves, two slits, two reflecting surfaces, and in every case a pattern of bright and dark fringes that the eye could read out as a wavelength. Now, we are diving to a different operation that the same wave model performs with only one of either: a single slit, a single sharp edge, a single hair across a beam. The phenomenon is called **diffraction**, and it is the answer to the puzzle Kipanga handed us in the very first paragraph of this chapter.

Recall the CD that Kipanga had brought to school on that long morning, the sharp white shaft of sunlight falling through the louvre windows, and the band of rainbow colour that walked across the silvery disc. There were no two slits, no thin film, no two reflecting surfaces. There was one shiny disc, one window full of morning light, and an outpouring of rainbow colour that geometric optics had no explanation for. Today we explain it. The grooves of a CD form a regular array of one thousand six hundred slits per millimetre, and each colour of the white sunlight bounces off that array at its own preferred angle. The CD is a diffraction grating, the most elegant single optical instrument the chapter will meet, and by the end of the section you will be able to read the wavelength of the sunlight's red component to two significant figures with nothing more than a CD case and a ruler.

The plan is two derivations and three protocols. The first derivation treats a single slit of width a, illuminated by a plane wave, and asks where the bright and dark bands fall on a distant screen. The second treats a regular array of many slits at common spacing d, the device called a diffraction grating, and asks the same question. Out of the two derivations come three formulas and one observation. The formulas read wavelengths out of angles. The observation tells you why a grating is preferable to a single slit, and why the CD that Kipanga held up to the bulb was, although he could not have known it on the day, doing physics that Newton's first telescope could not have done.

When diffraction shows itself

Diffraction is everywhere. Every time a wave passes an obstacle, the obstacle's edge sets off a chain of secondary wavelets that spread into the geometric shadow. The reason we do not normally see diffraction with light is a matter of scale. Visible light has wavelengths between about four hundred and seven hundred nanometres, which is to say between 4×10^{-7}m and 7×10^{-7}m. Everyday obstacles, like fingers, doors, and Mwanza chickens crossing a sunlit road, are between a centimetre and a metre wide. The ratio of obstacle width to wavelength is therefore between ten thousand and ten million, and the diffracted light spills only a few wavelengths into the shadow, which is to say a few micrometres on a length scale where the eye cannot resolve a hundred-micrometre fly. Diffraction is there. We simply cannot see it.

This is the concrete answer to a question that the chapter has been postponing. ***Why, despite light being a wave, does the everyday experience of light insist that it travels in straight lines?*** *Because everyday obstacles are very many wavelengths wider than the wavelengths themselves.* The geometric shadow is sharp to a precision set by the obstacle width divided by the wavelength, and that precision is in the millions for ordinary objects. Diffraction is not absent from your living room; it is merely too tightly confined to be noticed.

To see diffraction, we need an obstacle whose smallest dimension is comparable with the wavelength. A slit narrower than a hair, a wire thinner than a hair, a CD's groove pitch, an opal's array of microspheres. Once we get down to those scales, diffraction takes over and dominates. The rest of the section is about two such obstacles.

One more orientation, then to the algebra. A plane wave (a wavefront like a flat sheet of paper) hitting a slit produces a diffraction pattern whose shape depends on the geometry of where the screen sits. *If the screen is far enough away that the secondary wavelets reaching it can be treated as parallel,* we are in the regime called **Fraunhofer diffraction.** *If the screen is so close that the secondary wavelets reach it as diverging spherical waves,* we are in the regime called **Fresnel diffraction.** In school laboratories and in advanced-physics syllabuses we work exclusively in the **Fraunhofer regime;** the rest of the section assumes the screen is comfortably distant or, equivalently, that a converging lens has been used to bring parallel rays to focus.

The single slit

Take a single slit of width a, illuminated by a plane wave of wavelength λ arriving along the optical axis. The geometry is shown in the figure below. Far away on a screen at perpendicular distance L from the slit, we want to find the directions θ along which the diffracted intensity vanishes (the dark fringes) and the directions along which it is maximal (the bright fringes).

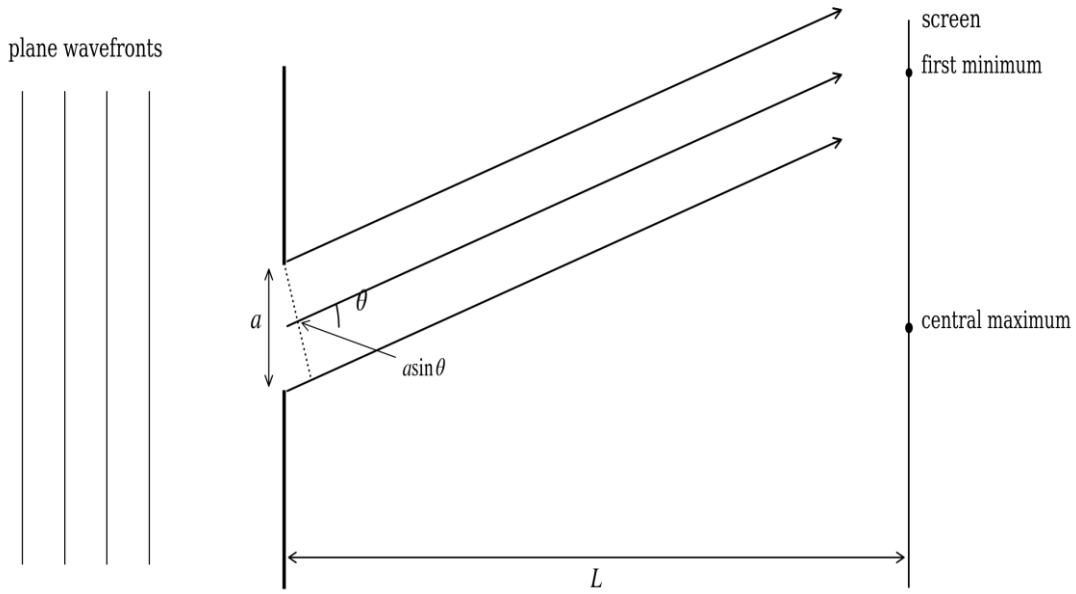


Figure: ***A plane wave arrives at a slit of width a.*** *Three representative rays leave the slit along the same angle θ to the optical axis. The path difference between the ray from the top of the slit and the ray from the bottom is asin θ. The path difference between any strip in the upper half and its partner directly opposite in the lower half is (a/2)sin θ. The first minimum occurs when this paired difference equals λ/2.*

The slit is not two sources but a continuous infinity of them. Every point on the wavefront crossing the slit acts as a secondary source (Huygens' principle, installed earlier in this chapter). To find the resultant at angle θ, divide the slit mentally into $2N$ equal strips, each of width $\frac{a}{2N}$, and pair the strip at position y in the upper half with the strip directly opposite, at position $y - a/2$, in the lower half. The path difference between any such paired pair, as their light leaves the slit at angle θ, is $(a/2)\sin\theta$, which is the horizontal distance $a/2$ projected on to the direction of propagation.

If this paired path difference equals exactly half a wavelength, the two strips arrive perfectly out of phase and cancel each other on the screen. Pair after pair cancels in the same way, and the whole slit contributes nothing at that angle. The condition for the first minimum is therefore:

$$\frac{a}{2}\sin\theta = \frac{\lambda}{2}$$

or, multiplying both sides by two,

$$\mathbf{a\sin\theta = \lambda}$$

The argument that produced the first minimum can be repeated by dividing the slit into four halves, six thirds, eight quarters, and so on. Each subdivision yields one more cancellation condition, and the general condition for a dark fringe in single-slit diffraction is:

$$\mathbf{a\sin\theta = \pm n\lambda, \quad n = 1, 2, 3, \ldots}$$

The $n = 0$ case is excluded because it corresponds to $\sin\theta = 0$, which is the optical axis itself, and along the optical axis every strip of the slit arrives in phase. This direction is the bright central maximum, and the diffraction pattern is symmetric about it: dark fringes appear in pairs at the same angle to either side, brought to the same position on the screen as $\mathbf{y = \pm n\lambda L / a}$ under the small-angle approximation $\sin\theta \approx \tan\theta \approx y / L$.

The intensity envelope and the width of the central maximum

Between the dark fringes lie bright bands, but they are not all the same brightness. A careful calculation of the resultant amplitude as a function of angle gives the intensity envelope shown in the next figure. The mathematical form is:

$$I(\theta) = I_0 \left(\frac{\sin\alpha}{\alpha}\right)^2, \quad \alpha = \frac{\pi a \sin\theta}{\lambda}$$

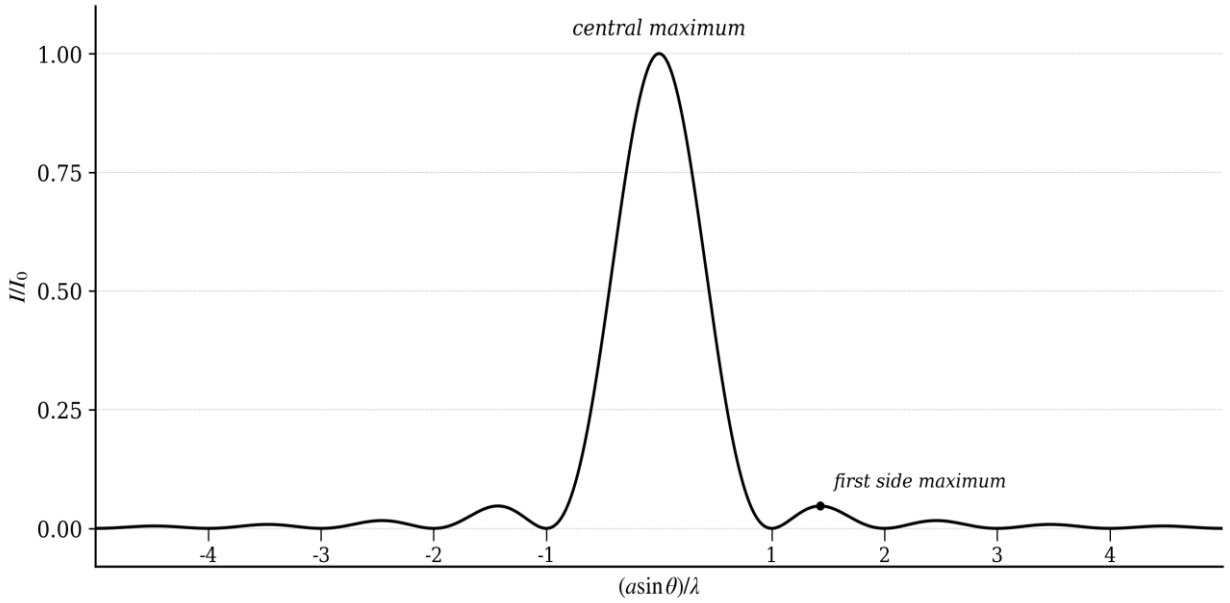


Figure: ***Intensity envelope I(θ) = I₀(sin α /α)² of single-slit diffraction, plotted as a function of (asin θ) /λ.*** *The central peak occupies the region from −1 to +1, the minima are at integer values, and the first side maximum carries only about 4.7% of the central peak intensity.*

Three features of the envelope deserve attention.

First, *the central maximum is twice as wide as any of the side maxima,* because the first minima occur at $\sin\theta = \pm\lambda/a$ while consecutive side minima are spaced by λ/a only. On a screen at distance L the full width of the central maximum is therefore:

$$W = \frac{2\lambda L}{a}$$

Second, *the side maxima are dim.* The first side maximum carries only about four-and-three-quarters per cent of the central peak intensity; the second side maximum is dimmer still. The vast majority of the diffracted light sits inside the central peak, and the rest is a faint corona that the eye notices only on a dark screen.

Third, *the position of every minimum and every maximum depends on the slit width a and the wavelength λ through the single ratio λ/a.* A narrower slit gives a wider pattern; a wider slit gives a narrower pattern. In the limit of a very wide slit the diffraction pattern collapses to a single bright spot on the optical axis, which is the geometric-optics result you have been trusting since O-level. Diffraction recovers geometric optics whenever the slit is many wavelengths wide.

The diffraction grating

Now repeat the experiment with not one slit but many. A **diffraction grating** *is a regular array of identical slits at a common spacing d* (the slit-to-slit distance, measured between the centres of adjacent slits). A school grating typically carries between three hundred and one thousand slits per millimetre, scratched into a thin glass plate so that the gaps act as the slits and the scratches act as the opaque material between them. The following figure shows the geometry.

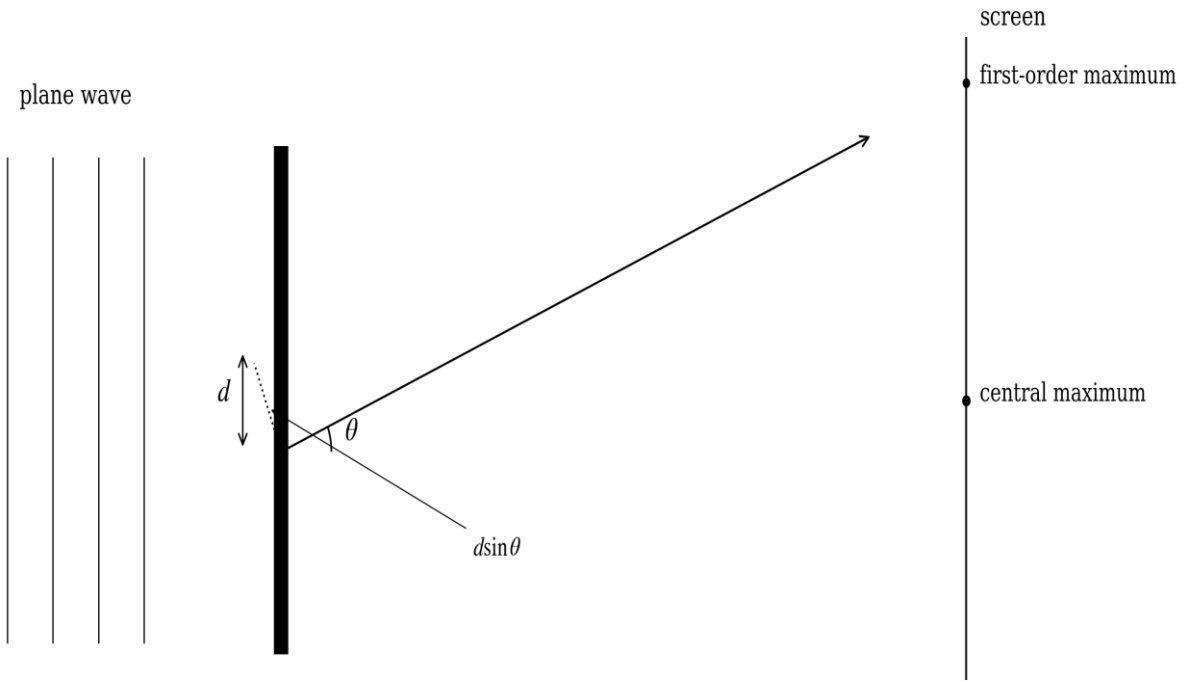


Figure: ***Plane wave incident on a grating of slit spacing d.*** *Two representative rays leave adjacent slits at the same angle θ to the optical axis. The path difference between them is dsin θ. The grating produces a bright principal maximum whenever this path difference equals an integer number of wavelengths.*

Apply Huygens' principle slit by slit. Each slit acts as a coherent secondary source, and the path difference between waves leaving two adjacent slits, at the common angle θ to the optical axis, is dsin θ. For all the slits to reinforce one another at that angle, the path difference must equal an integer number of wavelengths,

$$\mathbf{d\sin\theta = n\lambda}, \quad \mathbf{n = 0, \pm 1, \pm 2, \pm 3}, \ldots.$$

The integer n is called the order of the principal maximum. The $n = 0$ case is the central maximum at $\theta = 0$; the first-order pair sits at $\sin\theta = \pm\lambda/d$; the second-order pair at $\sin\theta = \pm 2\lambda/d$; and so on.

Two further conveniences of the grating formula are worth pinning down. The slit spacing d is most easily quoted as its reciprocal, the number of lines per unit length,

$$\mathbf{N = \frac{1}{d}},$$

so a grating with five hundred lines per millimetre has $d = (1/500)\text{mm} = 2.00 \times 10^{-6}\text{m}$. And the formula $d \sin\theta = n\lambda$ relates four quantities, any three of which determine the fourth. Knowing the wavelength and the spacing, one computes the angle at every order; knowing the angle and the spacing, one computes the wavelength of the source; knowing the angle and the wavelength, one computes the spacing of an unknown grating. The next two subsections exercise each pairing.

How many orders can be seen

Not every value of n produces a physical principal maximum. The grating equation requires $\sin\theta = n\lambda/d$, and the sine of a real angle cannot exceed unity. The highest order that the grating can produce on a screen is therefore the largest integer n satisfying $n\lambda \leq d$, or equivalently:

$$n \leq \frac{d}{\lambda}, \text{and thus:}$$

$$n_{max} = \text{floor}\left(\frac{d}{\lambda}\right)$$

The **floor** function (*round-down-to-the-nearest-integer*) is doing essential physical work here, not pedantry. There is no fractional principal maximum sitting on the screen at the angle implied by a non-integer ratio d/λ. The next dialogue makes the point precisely.

Kipute: *Sir, for a grating with five hundred lines per millimetre, illuminated by green light of wavelength 526nm, the ratio d divided by lambda comes to 3.80. So the maximum order is four, after rounding up.*

Mr. Akilikubwa: *Stop, Kipute. Read the formula once more. n_{max} is the floor of d/λ, not the rounded value. Why floor?*

Kipute: *Because sin theta cannot exceed one.*

Mr. Akilikubwa: *Correct. So whatever n_{max} is, it is the largest integer for which sin θ comes out at less than or at most equal to one. For your grating, n equals three gives sin θ equal to 3 times 526 nanometres divided by 2 micrometres, which is 0.789. That fits. n equals four would need sin θ equal to 1.052, which does not fit; an angle whose sine is bigger than one does not exist. So the answer is three. There is no 3.8th maximum sitting on the screen. There is the third maximum at one angle, and after that, geometry has no further angles to offer.*

Kipute: *Always floor.*

Mr. Akilikubwa: *Always. The floor function in this formula is not a mathematical convention; it is the geometric statement that sines top out at one. The day a sine exceeds one is the day a new physics is being discovered.*

Why a grating is sharper than a double slit

Both Young's two-slit setup and the diffraction grating obey the same constructive-interference condition $d\sin\theta = n\lambda$. The angles at which the bright fringes appear are therefore identical. *Yet the grating's lines are visibly sharper, narrower, and brighter than Young's. Why?*

The answer is in the number of slits, N. With two slits the resultant amplitude oscillates between full reinforcement (every $d\sin\theta = n\lambda$) and zero (every $d\sin\theta = (n + ½)\lambda$) as the angle is swept. With N slits, full reinforcement still occurs at the same angles, but the cancellation between them is far more thorough; between consecutive principal maxima the resultant amplitude oscillates many times and passes through $N - 1$ subsidiary minima and $N - 2$ much weaker subsidiary maxima. The principal maxima themselves become narrower in inverse proportion to N:

$$\Delta\theta_{\text{principal}} \propto \frac{1}{N d\cos\theta}$$

A grating with a thousand slits illuminated across its width therefore gives principal maxima a thousand times narrower than a two-slit pattern at the same d and λ. This is the property that makes the diffraction grating, and not Young's two slits, the tool of choice for measuring wavelengths precisely. A double-slit pattern can locate a wavelength to perhaps one per cent; a grating can locate the same wavelength to one part in a million. The principle is identical; the resolution differs by six orders of magnitude.

Two wavelengths sharing a direction

A light source rarely emits a single wavelength. A sodium lamp emits two closely spaced yellow lines at 589.0nm and 589.6nm; a mercury vapour lamp emits a spectrum of distinct lines; daylight emits a continuum. When such a source illuminates a grating, each wavelength produces its own family of principal maxima, and the families intersect at angles where two wavelengths happen to share a common diffraction angle. The condition for such a coincidence is simply

$$\mathbf{n_1\,\lambda_1 = n_2\,\lambda_2} = d\sin\theta,$$

with n_1 and n_2 positive integers and the same product $d\sin\theta$ on both sides. Given the two wavelengths, the smallest positive integers (n_1, n_2) satisfying this equation locate the smallest-angle coincidence; the second coincidence sits at $(2n_1, 2n_2)$, and so on. Several examination problems set the integers and the wavelengths and ask for the angle; several others set the integers and the angle and ask for the missing wavelength. The arithmetic is the same in both directions.

Measuring a wavelength with the spectrometer

All of the above is at its most useful in the school's spectrometer, a compact instrument designed exactly to measure the angles between the principal maxima of a grating, and from them the wavelengths of a source. Figure below shows the layout.

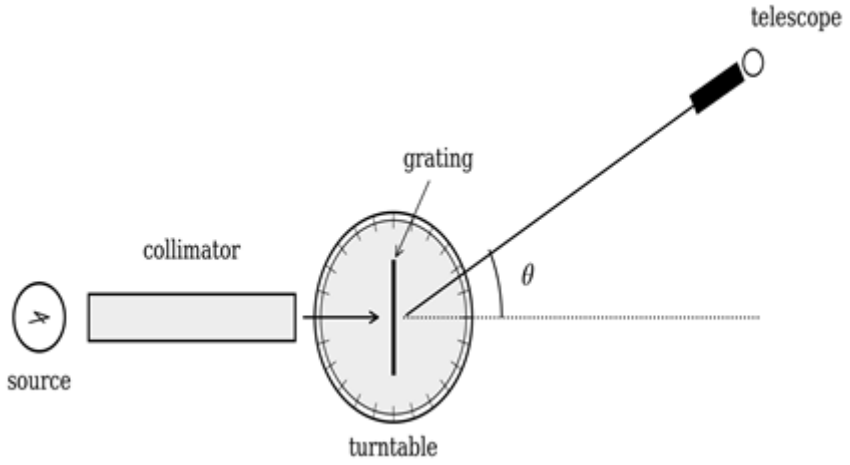


Figure: ***Bench-top spectrometer arrangement.*** *The source is imaged through a fine slit by the collimator; the parallel beam falls on a grating mounted on the turntable; the diffracted parallel rays are collected by the rotatable telescope, whose pointing angle is read off the turntable scale.*

The instrument has three optical components and one calibrated angular scale. The collimator turns the diverging light from the source slit into a parallel beam. The grating, mounted on a rotatable turntable with its plane perpendicular to the incident beam, diffracts the parallel beam into its principal maxima. The telescope, also rotatable on the same turntable, is swung along the angular scale until its cross-wires sit on a chosen principal maximum, and the angle read off the scale is recorded.

Because the turntable's zero is rarely exactly at $\theta = 0$, the standard practice is to swing the telescope to the principal maximum on the right of the central maximum (recording angle θ_R), then swing it across centre to the corresponding principal maximum on the left (recording angle θ_L), and take the half-difference,

$$\theta = \frac{\theta_R - \theta_L}{2}$$

This neatly removes whatever zero-error the turntable carries. With θ so measured and d known from the grating's specification, the wavelength of the source is computed from the grating equation,

$$\lambda = \frac{d\sin\theta}{n}$$

Repeating the measurement at several orders and averaging is the standard refinement. With a thousand-lines-per-millimetre grating, a school spectrometer routinely measures sodium yellow to within one nanometre, which is the precision that decided the discovery of helium on the sun in 1868 before any chemist had isolated it on Earth.

Five worked examples now follow. They are where the formulae stop being decoration and start being instruments.

BINDER Example 17

A laser of wavelength 650nm illuminates a narrow slit. A diffraction pattern is observed on a screen at perpendicular distance 4.0m behind the slit. The two dark fringes that bracket the central bright band are measured to be 52mm apart. Calculate the width of the slit.

Solution

Let a be the slit width, $\lambda = 6.50 \times 10^{-7}\text{m}$ the wavelength, $L = 4.0\text{m}$ the slit-to-screen distance, and $W = 52 \times 10^{-3}\text{m}$ the full width of the central maximum. The full width is the distance between the first minima on either side of the centre, namely:

$$W = \frac{2\lambda L}{a}$$

Rearrange for a:

$$a = \frac{2\lambda L}{W}$$

Insert numerical values:

$$a = \frac{2 \times (6.50 \times 10^{-7}\text{m}) \times (4.0\text{m})}{52 \times 10^{-3}\text{m}} = 1.00 \times 10^{-4}\text{m},$$

which is 0.100mm, or $100\mu\text{m}$. The slit is therefore one tenth of a millimetre wide.

Making Sense of the Answer: *A slit width of 100 micrometres is right in the range that produces visible diffraction with red laser light. Anything much wider would give a central maximum thinner than the laser dot itself, and the diffraction would not be visible to the eye; anything much narrower would spread the central maximum so wide that the dark fringes would leave the screen. The fact that the answer comes out at exactly the order of magnitude a school laser slit was designed for is a strong sanity check on the arithmetic.*

Think Like a Physicist: *Whenever a single-slit problem gives you the full width of the central maximum, write $W = 2\lambda L/a$ and rearrange. Whenever it gives the distance from the centre to the first minimum, write $y_1 = \lambda L/a$ and rearrange. The two formulas are the same formula, differing only by a factor of two for the full-width-versus-half-width convention. Read the problem carefully to know which one you have been handed.*

BINDER Example 18

A diffraction grating with 667 lines per millimetre is illuminated normally by monochromatic light of wavelength 530nm. Calculate the angles to the normal at which the first-order and second-order principal maxima are observed.

Solution

Let $N = 667\ \text{lines/mm} = 6.67 \times 10^5\ \text{lines/m}$, so the slit spacing is :

$$d = \frac{1}{N} = \frac{1}{6.67 \times 10^5\ \text{m}^{-1}} = 1.50 \times 10^{-6}\text{m}.$$

The grating equation reads $d\sin\theta = n\lambda$, which rearranges to:

$$\sin\theta = \frac{n\lambda}{d}$$

First-order ($n = 1$):

$$\sin\theta_1 = \frac{1 \times 5.30 \times 10^{-7}\text{m}}{1.50 \times 10^{-6}\text{m}} = 0.3533,$$
$$\theta_1 = 20.7°.$$

Second-order ($n = 2$):

$$\sin\theta_2 = \frac{2 \times 5.30 \times 10^{-7}\text{m}}{1.50 \times 10^{-6}\text{m}} = 0.7067,$$
$$\theta_2 = 45.0°.$$

The first-order principal maximum sits at $20.7°$ and the second-order at $45.0°$ to the normal.

Making Sense of the Answer: *The second-order angle is a little more than twice the first-order angle, not exactly twice. The doubling of sin θ does not produce a doubling of θ because sine is not a linear function. This non-linearity is the reason a grating's higher orders sit closer together as the order increases, and is the reason gratings are easier to read in the first two or three orders than in the eighth or ninth. The arithmetic also exposes how easy it is to misread the grating equation: students who confuse sine and tangent here will get the right answer at small angles and very wrong answers at large ones. Stick to sine.*

Think Like a Physicist: *In every grating problem, convert lines per millimetre to slit spacing in metres before you do anything else. Carrying the factor of a thousand through the rest of the algebra is the most common source of an answer that is wrong by exactly that thousand-fold.*

REAL Example 19

A spectrometer is set up to examine yellow sodium light of wavelength 580nm using a diffraction grating ruled with 500 lines per millimetre. Calculate the highest order of principal maximum that can be observed and, separately, the angle at which it appears.

Solution

The slit spacing is:

$$d = \frac{1}{5.00 \times 10^5\ \text{m}^{-1}} = 2.00 \times 10^{-6}\text{m}$$

Sines have an upper limit of one. The grating equation rearranged for the order is

$$n = \frac{d\sin\theta}{\lambda},$$

so the largest integer n that the grating can accept must satisfy $n\lambda \leq d$, or

$$n_{max} = \text{floor}\left(\frac{d}{\lambda}\right) = \text{floor}\left(\frac{2.00 \times 10^{-6}\text{m}}{5.80 \times 10^{-7}\text{m}}\right) = \text{floor}(3.448) = 3.$$

The highest visible order is $n = 3$. Its angle follows from the grating equation:

$$\sin\theta_3 = \frac{3 \times 5.80 \times 10^{-7}\text{m}}{2.00 \times 10^{-6}\text{m}} = 0.870,$$
$$\theta_3 = 60.5°.$$

The third-order principal maximum is the highest visible. It appears at an angle of $60.5°$ to the normal on each side of the central maximum.

Making Sense of the Answer: *The floor function is doing the work of keeping the geometry honest. A naive rearrangement of d/λ gives 3.45, and the temptation is to round that up to four. But the fourth order would require sin θ to be greater than one, which is geometrically impossible. The grating quietly refuses to deliver a fourth maximum, and the floor function is the mathematical statement of that refusal. The remaining third-order maximum at 60.5 degrees is consistent: sin 60.5 degrees is 0.870, which matches the algebra, and a beam diffracted at 60 degrees is a beam diffracted at a very large angle, so the third order is sitting close to the grating's geometric limit.*

Think Like a Physicist: *When a problem asks for the highest order, compute d divided by lambda, then immediately apply the floor function before doing anything else. Resist the temptation to round in either direction. Only after the floor function gives an integer should you compute the angle, and that angle should always come out at less than ninety degrees if the floor was applied correctly.*

HOT Example 20

In a spectrometer experiment, a diffraction grating illuminated normally by blue light of wavelength 480nm produces its second-order principal maximum at an angle of $35.2°$ to the normal. Calculate the number of lines per millimetre of the grating.

Solution

Let N be the number of lines per metre and $d = 1/N$ the slit spacing. The grating equation rearranges for the spacing as:

$$d = \frac{n\lambda}{\sin\theta}$$

Substituting numerical values (using $n = 2$ and $\sin 35.2° = 0.5764$):

$$d = \frac{2 \times 4.80 \times 10^{-7}\text{m}}{0.5764} = 1.665 \times 10^{-6}\text{m}$$

Therefore:

$$N = \frac{1}{d} = \frac{1}{1.665 \times 10^{-6}\text{m}} = 6.00 \times 10^{5}\text{m}^{-1}$$

Converting to lines per millimetre:

$$N = \frac{6.00 \times 10^{5}}{10^{3}\text{mm}} = 600\ \text{lines/mm}$$

The grating has 600 lines per millimetre.

Making Sense of the Answer: *600 lines per millimetre is bang in the middle of the laboratory range for ruled diffraction gratings, which typically run from 100 lines per millimetre at the coarse end (used for the infrared) up to 2400 lines per millimetre at the fine end (used for the ultraviolet). A blue line at 480 nanometres diffracting at 35 degrees in the second order is a sensible setup that anyone with a school spectrometer can reproduce. Had the answer come out at 60000 lines per millimetre or at 60 lines per millimetre, we would suspect a factor-of-a-thousand slip in the arithmetic; coming out at 600 is the right neighbourhood.*

Think Like a Physicist: *The grating equation always carries four quantities: order, wavelength, angle, and spacing. Any three of the four determine the fourth. Reading the problem to identify which three are given is more than half the work. Once you know the three you have, write the equation with the fourth quantity on its own on the left, and the substitution is mechanical from there.*

HOT Example 21

A diffraction grating of slit spacing 2.0μm is illuminated normally by light containing two wavelengths. The third-order principal maximum of the longer wavelength is observed to fall on the screen at exactly the same angle as the fourth-order principal maximum of the shorter wavelength. The longer wavelength is 600nm. Calculate the shorter wavelength, and the angle at which the two maxima coincide.

Solution

Let $\lambda_1 = 6.00 \times 10^{-7}\text{m}$ be the longer wavelength and λ_2 the shorter, and let θ be the common angle at which the two principal maxima coincide. The two grating equations are $d\sin\theta = 3\lambda_1$ and $d\sin\theta = 4\lambda_2$. Equating the two right-hand sides,

$$3\lambda_1 = 4\lambda_2, \quad \text{so} \quad \lambda_2 = \frac{3\lambda_1}{4}$$

Substituting numerical values:

$$\lambda_2 = \frac{3 \times 6.00 \times 10^{-7}\text{m}}{4} = 4.50 \times 10^{-7}\text{m} = 450\text{nm}$$

The shorter wavelength is 450nm. Its coincidence angle follows from either of the two grating equations:

$$\sin\theta = \frac{3\,\lambda_1}{d} = \frac{3 \times 6.00 \times 10^{-7}\text{m}}{2.00 \times 10^{-6}\text{m}} = 0.900,$$
$$\theta = 64.2°.$$

The two maxima coincide at 64.2° to the normal.

Making Sense of the Answer: *The two wavelengths span the orange end of the spectrum (600nm) and the blue end (450nm), which together form a fair share of the visible band. The fact that two such different colours can share an angle on the same grating is one of the reasons why high-order spectra are difficult to interpret without prior knowledge of which colours are present; the second-order red can sit on top of the third-order green, the third-order orange can sit on top of the fourth-order blue, and so on. Spectrometers reserve their highest-precision work for the lower orders where coincidences are rarer.*

Think Like a Physicist: *Whenever a problem mentions that two orders of two wavelengths share an angle, set the two right-hand sides of the grating equation equal to each other and solve for the missing wavelength. The slit spacing drops out of that step; it returns only when you compute the actual angle. Two parts, both clean, neither difficult.*

We have now extracted, from a single slit and from a grating, three formulas that connect the four quantities of any diffraction experiment: the slit width or spacing, the wavelength of the source, the angle of the principal maximum, and the order. With these in hand, the chapter has met its second physical question. A wave passing one obstacle of size comparable with the wavelength produces a pattern of intensity whose geometry is completely determined by the obstacle's size and the wave's wavelength.

Two physical operations now have a complete account: superposition of two beams (interference) and superposition within a single beam passing an obstacle (diffraction). In a real apparatus, both happen at once. A Young's double-slit pattern is not, strictly speaking, made only of two-source interference fringes; it is made of those fringes modulated by the diffraction envelope of each single slit. A grating pattern is not, strictly speaking, only the constructive sum of N sources; the diffraction envelope of each slit weights the principal maxima by an angle-dependent factor. The next short reconciliation makes the relationship between the two operations explicit and answers, in six clean differences.

INTERFERENCE VERSUS DIFFRACTION

The chapter has spent its pages on two physical operations, and the time has come to set them side by side and admit what is the same and what is different. **Interference** asks *what two coherent waves do when they arrive at the same point.* **Diffraction** asks *what one wavefront does when it passes an obstacle whose size is comparable with the wavelength.* The two questions sound different. Their answers, on closer inspection, are made of the same ingredient: **superposition.** Where they part company is in what is superposing, how many of it there are, and whether the superposing pieces were physically separated to begin with.

A six-row comparison brings the differences to the surface in the cleanest available form.

Property	**Interference (two-source)**	**Diffraction (single-source spreading)**
Source structure	Two coherent sources of comparable amplitude, physically separated.	One wavefront passing one obstacle of size comparable with λ.
Origin of the pattern	Superposition of two whole waves.	Superposition of infinitely many secondary wavelets from a single obstacle.
Fringe widths	All fringes are equally spaced.	Central maximum is twice as wide as any of the side maxima.
Intensity of fringes	All bright fringes have the same intensity.	Intensity falls off rapidly with order; the envelope is $(\sin\alpha/\alpha)^2$.
Order count	Bounded by $\sin\theta \leq 1$: $n_{max} \approx d/\lambda$; large for typical d.	Bounded by $\sin\theta \leq 1$: $n_{max} \approx a/\lambda$; smaller, because a single slit is narrower than a slit spacing.
Visibility	Crisp fringes visible without lenses, simply by projecting the screen far behind the slits.	Often requires a converging lens to focus the diffracted parallel rays onto a finite screen.

The two columns above are the differences. What follows is the unity behind them: in every real apparatus, both operations happen at once, and the pattern the eye finally sees is the product of the two.

Why both phenomena coexist in every real apparatus

Take Young's double-slit experiment, the first apparatus the chapter built. The setup is two narrow slits, separated by a distance d, illuminated by a coherent wavefront. The interference formula tells us that the screen will carry bright fringes at every angle where the path difference between the two slits is an integer number of wavelengths,

$$d\sin\theta = n\lambda$$

That is the truth, but not the whole truth. A real slit is not a mathematical line. It is a slit of some finite width a, and the wave passing through it does not emerge as a clean spherical wavelet; it emerges as a full single-slit diffraction pattern with the intensity envelope we met in the previous section,

$$I_{diff}(\theta) \propto \left(\frac{\sin\alpha}{\alpha}\right)^2, \quad \alpha = \frac{\pi a \sin\theta}{\lambda}$$

The two-slit interference pattern, with peaks at $d\sin\theta = n\lambda$, is therefore multiplied at every angle by the single-slit envelope. What appears on the screen, $I_{obs}(\theta)$, is the product:

$$I_{obs}(\theta) \propto \cos^2\left(\frac{\pi d \sin\theta}{\lambda}\right) \times \left(\frac{\sin\alpha}{\alpha}\right)^2$$

Three observations explain why this product is the right picture of what is actually seen at the school bench.

First, *the cos² factor oscillates much faster than the (sinα/α)² envelope,* because the ratio d/a is typically four or five for a Young's-slit card. The envelope wears the cos²-fringes as a slow modulation. **Second,** *the angle θ at which the diffraction envelope reaches its first zero, $a\sin\theta = \lambda$, is far larger than the angle θ at which the interference reaches its first zero, $d\sin\theta = \lambda/2$,* so most of the visible fringes sit inside the central diffraction lobe and the eye sees a few orders of bright cos² fringes before the envelope dims them into invisibility. **Third,** *the fringe whose centre would have been at the diffraction minimum is missing from the screen entirely,* because the envelope multiplies it by zero. This is the so-called **missing order**, and it gives a quick experimental measure of the ratio d/a: count the number of bright cos² peaks inside the central diffraction lobe, divide by two, and the result is d/a almost exactly.

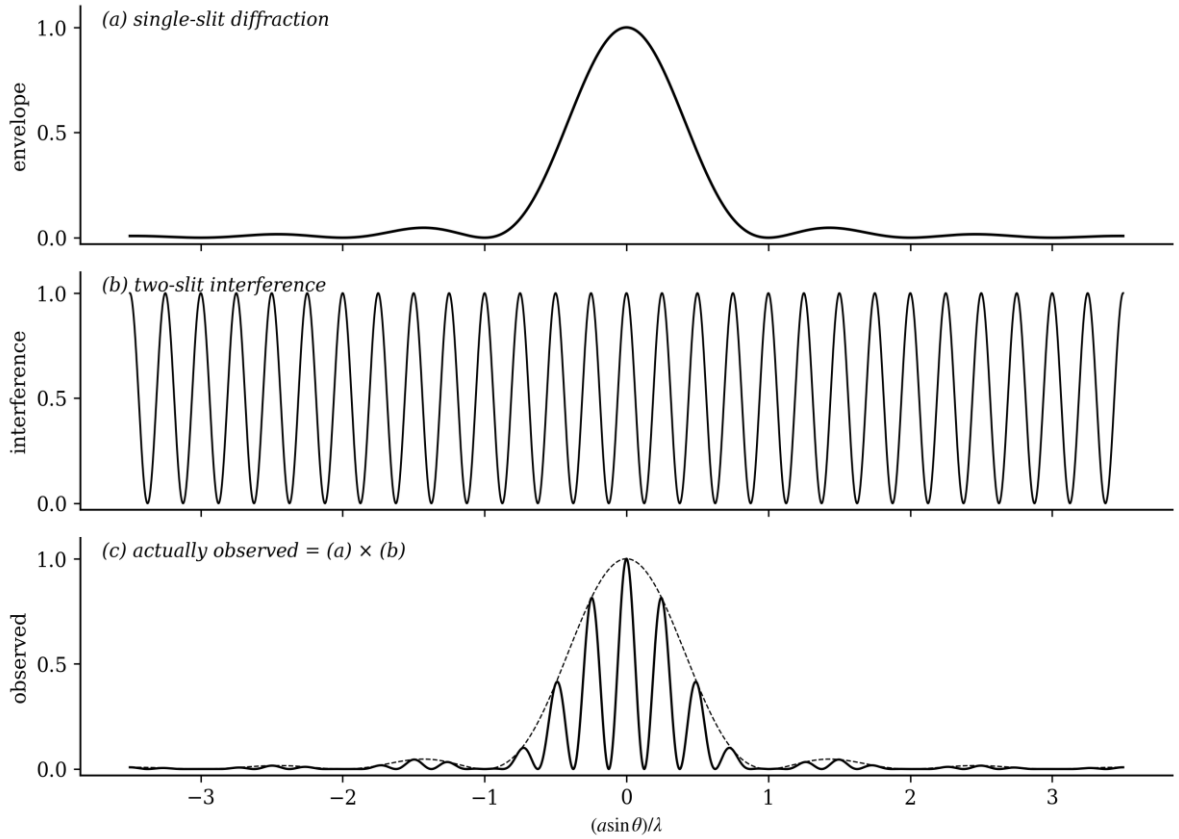


Figure: ***How the two operations combine on the screen in Young's double-slit experiment****. Panel (a) is the diffraction envelope of one slit on its own. Panel (b) is the interference pattern of two infinitesimally narrow slits on their own. Panel (c), the pattern an experimenter actually observes, is the point-by-point product of the two. The fast oscillation in (c) is interference; the slow shape that wraps around it is diffraction; the small fringes that go missing where the envelope passes through zero are the missing orders.*

Two further points generalise the picture to a diffraction grating. The grating with N slits replaces the cos² interference factor with the much sharper N-slit interference function, whose principal maxima are at the same angles $d\sin\theta = n\lambda$ but whose linewidth shrinks in inverse proportion to N. The diffraction envelope of each individual slit is unchanged. The actually observed pattern is, once again, the product of an interference factor that depends only on d and an envelope that depends only on a. The envelope is the slower; the interference is the faster; and the missing orders of Young's slits become, in the grating, suppressed orders that the manufacturer routinely chooses where to place by tuning d/a during the ruling.

So the difference between interference and diffraction is a difference of vocabulary. Once both operations are admitted to the same apparatus, the difference between them is no longer a contradiction; it is a multiplication. The wave model that has carried the chapter for about seventy pages does not care whether the screen pattern is read out as fringes inside an envelope or as the product of two superposition sums. The two readings are the same thing said twice.

Two questions of the chapter have now received their full account. Where the wave goes when two waves arrive at one point, and where the wave goes when one wave passes one obstacle, are both, in the end, exercises in adding amplitudes and squaring. A third question remains for the chapter to answer. The wave has been described, all the way through, as a transverse disturbance. The direction of the disturbance, the way the wave shakes as it travels, has been treated as a quiet detail in the

background. It is time to bring that detail forward. The next section is about polarisation, the way in which a transverse wave commits itself to a single plane of vibration and what happens when you ask it to commit to a different one.

POLARISATION

The chapter has so far asked the wave only two questions. *Where does it go when two of itself meet at a point? Where does it go when one of itself meets a slit?* Both questions take the wave as a scalar quantity, a height above the average that adds, subtracts, and squares. There is a third question that the previous two have left unattended. The wave is transverse: it shakes across the direction in which it travels. Across is not one direction. It is two, and within those two there is an infinite family of intermediate directions. The question of polarisation is the question of which of those directions the wave has chosen, and what happens when an apparatus or a medium asks it to choose differently.

Two reasons make this a topic the syllabus cannot afford to bury at the back. First, polarisation is the simplest experimental signature by which a transverse wave is told apart from a longitudinal wave. Sound, the most familiar wave of school physics, is longitudinal and cannot be polarised; light, the wave the chapter has been studying, is transverse and can. Second, polarisation has wandered out of the laboratory and into the engineering of nearly every optical device the candidate will hold during their working life. Liquid-crystal screens, photographic filters, three-dimensional cinema spectacles, and the strain test of a windscreen all depend on it. The opening pages of this section build the idea from a single experiment, derive the law that governs it, and then take a quick tour of the four ways in which plane-polarised light can be produced.

A test that tells a transverse wave from a longitudinal one

Take two identical pieces of Polaroid film, the kind sold for camera filters or polarising sunglasses. Place one in front of a window through which the morning sun is falling on the louvre slats of the Miono Secondary classroom. Looking through the single film, the world is dimmed but still visible. Now place the second film in front of the first and slowly rotate it about the line of sight. The brightness rises and falls; at one orientation the view is at its brightest, at the orientation a quarter-turn away it has gone almost completely dark, and the cycle repeats every half-turn. Two crossed Polaroids extinguish the light.

This is not what would happen if light were a longitudinal wave. A longitudinal disturbance shakes along the direction of travel, and there is only one such direction; nothing rotated about the line of travel could ever block such a wave. The fact that two films can be arranged to extinguish the light, and then unblock it by a quarter-turn, proves at once that whatever is shaking inside the wave is shaking sideways to its motion, and that the sideways direction can take more than one orientation. The wave is transverse, and the two films are sorting it by the orientation of its shake. The film that emits the sorted wave is the **polariser**; the second film, used to detect the orientation that the first has chosen, is the **analyser**.

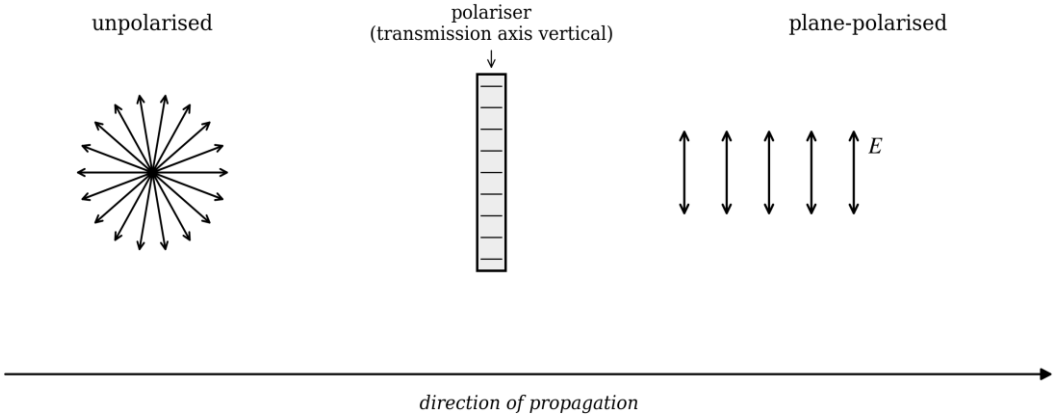


Figure: *Unpolarised light has its transverse vibrations spread evenly over every direction perpendicular to the propagation direction; on the left of the diagram, each short double-headed arrow stands for one such oscillation direction, picked at random. A polariser passes only the component along its transmission axis, which here is drawn vertical; what emerges on the right is plane-polarised light, every oscillation now confined to a single direction.*

Plane of vibration and plane of polarisation

Two pieces of geometric language need fixing before any law is written down. A light wave is, as the candidate already knows from earlier chapters, a coupled pair of oscillating fields: an electric field, conventionally written **E** and pointing along the direction in which the wave shakes its charges, and a magnetic field, conventionally written **B** and pointing along a direction perpendicular to E and also perpendicular to the line of travel. The two fields are tied together: wherever E is large, B is also large; wherever E reverses sign, so does B. The three directions, E, B, and the direction of travel, form a right-angled triple.

With E identified as the line in which the wave actually shakes, the geometry is straightforward. The plane that contains the line of E and the direction of travel is the **plane of vibration**. The plane that is perpendicular to it, and which also contains the direction of travel, is the **plane of polarisation**; this is the plane in which B lies. The two planes intersect along the

direction of travel and are at right angles to each other. Older textbooks sometimes use the two names the other way round, so the candidate should be ready to identify which of the two planes is meant by inspecting whether the author is talking about E or about B. Throughout this book the convention is that the plane of vibration is the plane of E, because E is the field that polarisers, analysers, photographic plates, and biological detectors actually respond to.

Three states of polarisation will appear in what follows.

- **Unpolarised** light has its electric vector pointing in every direction perpendicular to the line of travel with equal probability; an ordinary lamp, the Sun, and a candle all emit such light.
- **Plane-polarised** (or **linearly polarised**) light has its electric vector confined to one fixed plane; the output of a Polaroid filter is of this kind.
- **Partially polarised** light has a preferred direction of vibration but not an exclusive one; light reflected off a window pane at most angles is of this kind.

The remainder of this section concerns plane-polarised light, because the syllabus and the laboratory both treat it as the case from which everything else follows.

Four ways to make plane-polarised light

Nature does not insist that plane-polarised light be produced by a Polaroid film. Four physically distinct mechanisms are catalogued by the syllabus and will be examined in turn in the next four subsections. They share no apparatus, but they share one outcome: every one of them converts unpolarised light into a wave whose electric vector lies in a single plane. The four mechanisms are:

1) **Selective absorption** by an oriented molecular grid (the Polaroid filter);
2) **Reflection** at a denser medium at a particular angle (Brewster's law);
3) **Double refraction** inside an anisotropic crystal followed by suppression of one of the two refracted beams (the Nicol prism); and
4) **Scattering** by particles smaller than the wavelength (the blue of the daytime sky).

The first two are the most often examined and will receive the longest treatment; the third introduces the most beautiful piece of nineteenth-century optical engineering ever placed on a school syllabus; the fourth explains why the sky is the colour it is.

Polaroid film and Malus's law

The film named after the company that first manufactured it consists of long-chain organic molecules, drawn during manufacture into parallel alignment and then doped with iodine atoms that make them electrically conducting along their length. The electric field of a light wave whose vibration is parallel to the molecular chains drives free charge up and down the chain, loses its energy to the chain as heat, and is therefore absorbed. The component whose vibration is perpendicular to the chains finds no charge to push, dissipates no energy, and is transmitted. The transmission axis of the Polaroid is, perhaps counter-intuitively, perpendicular to the direction in which the molecules lie.

Unpolarised light incident on such a film has its electric vector pointing equally in every direction. The component perpendicular to the molecular chains is half of the total energy on average; the component parallel to the chains, also half, is absorbed. A perfect Polaroid film therefore halves the intensity of unpolarised light and turns the transmitted half into plane-polarised light whose plane of vibration is the transmission axis of the film. Write this as a numerical statement to fix it: if I_{unp} is the intensity of unpolarised light, then after one Polaroid film,

$$\mathbf{I = \frac{1}{2} I_{unp}}$$

Now the geometry of two films. Pass plane-polarised light, of intensity I_0 and electric-vector amplitude E_0, onto a second Polaroid whose transmission axis makes an angle θ with the plane of vibration of the incoming wave. Decompose E_0 into one component along the analyser's transmission axis and one perpendicular to it. The component along the axis has magnitude $E_0\cos\theta$; the perpendicular component is absorbed. The transmitted amplitude is therefore $E = E_0\cos\theta$, and the transmitted intensity, being proportional to the square of the amplitude, is:

$$\mathbf{I = I_0 \cos^2 \theta}$$

This is **Malus's law**, named for Étienne-Louis Malus, who first wrote it down in 1809 after watching the setting Sun reflected in the windows of the Luxembourg Palace through a piece of Iceland spar. In words:

When plane-polarised light passes through an ideal polarising film, the intensity of the transmitted beam is equal to the intensity of the incident beam multiplied by the square of the cosine of the angle between the plane of vibration of the incident light and the transmission axis of the film.

The law will return in every numerical problem that involves two or more polarising films in series. Notice three details. At $\theta = 0$ the analyser passes the entire incident intensity, because $\cos^2 0 = 1$. At $\theta = \pi/2$, that is, with axes crossed, it passes nothing, because $\cos^2(\pi/2) = 0$; this is the extinction the morning observer saw when crossing the two films. At $\theta = \pi/4$ the transmitted intensity is one half of the incident intensity, because $\cos^2(\pi/4) = 1/2$.

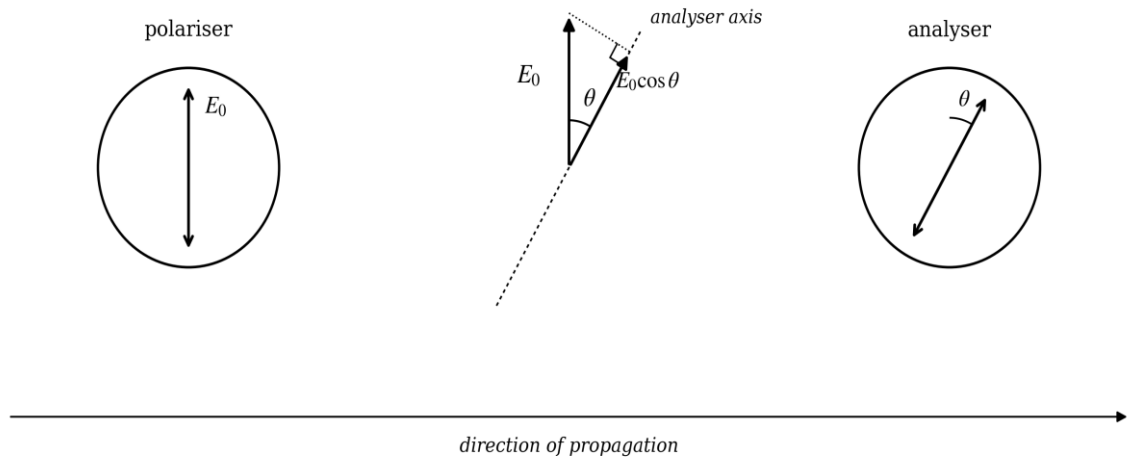


Figure: ***The geometry from which Malus's law follows.*** *Plane-polarised light leaves the polariser with electric vector along the vertical transmission axis, of magnitude E_0. The analyser's transmission axis makes an angle θ with the polariser's axis. The component of E_0 along the analyser axis is $E_0\cos\theta$; the transmitted intensity is therefore proportional to the square of this projected amplitude.*

Three polaroids in series

Malus's law has a consequence that is easy to state and unsettling to think about until it is grasped. Cross two polaroids so that no light comes through. Now slide a third polaroid between them at forty-five degrees to either of the outer two. The light that had been extinguished is reborn. A film placed in the path of an already extinguished beam apparently emits light it does not receive. The puzzle is sharp enough to merit a dialogue.

Kipanga: *Mr. Akilikubwa, two polaroids at 0° and 90° give zero light. I just put a third in the middle and the room got brighter behind the third one. The polaroid is making light?*

Mr. Akilikubwa: *The polaroid is not making any light. It is rearranging the light that does pass it. Step through it with me one stage at a time, please.*

Kipute: *First polaroid is vertical. Unpolarised light goes in; half of it comes out, plane-polarised vertical. Intensity $I_0/2$.*

Mr. Akilikubwa: *Correct. Now the middle polaroid is at 45° to the vertical. What does Malus's law say?*

Kipanga: *Transmitted intensity is the incoming intensity multiplied by* $\cos^2(45°)$*, which is 1/2. So we have $(I_0/2) \times (1/2) = I_0/4$ coming out of the middle polaroid, with its plane of vibration now along the middle polaroid's axis at 45°.*

Mr. Akilikubwa: *Good. Now the last polaroid is at 90° from the first, which is 45° from the middle one. Apply Malus's law again.*

Kipute: *$(I_0/4) \times \cos^2(45°) = I_0/8$. So we get one eighth of the original intensity, not zero.*

Mr. Akilikubwa: *And now compare with the case where the middle polaroid is taken out. Two polaroids at 0° and 90°: Malus's law gives $I_0/2$ multiplied by $\cos^2(90°)$ which is zero. Removing the middle polaroid extinguishes the light completely; putting it back gives back $I_0/8$.*

Kipanga: *So the middle one is not emitting light from itself; it is changing the plane of vibration so that something is allowed through the third one. A polaroid is not a curtain that blocks light. It is a rotor that turns the plane of vibration to a new direction and discards what cannot fit.*

Mr. Akilikubwa: *Exactly. A polaroid in the path of plane-polarised light absorbs the perpendicular component and keeps the parallel one. Without the middle film, the third film is being asked to pass light whose plane is at 90° to its own axis; nothing to project, nothing to keep. With the middle film, the third film is being asked to pass light whose plane is at 45° to its own axis; half by amplitude, a quarter by intensity, survives. The same physics in both cases.*

The dialogue makes a general point. Adding a polariser to the path of a plane-polarised beam can never increase its intensity; that is obvious from the $\cos^2\theta$ factor, which is at most one. Adding a polariser to the path between two crossed polarisers can however increase the intensity from zero to a positive value, because the second polariser, no longer required to absorb the parallel component of a beam, is being given a beam whose plane has been rotated. Polarisers are not curtains; they are decomposers, keeping the part of the wave that fits and absorbing the rest.

Brewster's law

A second route to plane-polarised light requires no special film. Allow unpolarised light to fall on the surface of glass, water, or any transparent dielectric medium. Some of the light reflects, some refracts. At most angles, both the reflected beam and the refracted beam are partially polarised; the reflected component has a slight excess of vibrations parallel to the surface and the refracted component has a slight excess perpendicular to the surface. At one special angle of incidence, however, the reflected beam is **completely** plane-polarised, with its electric vector lying parallel to the reflecting surface and perpendicular to the plane of incidence. That special angle is the **polarising angle** of the medium and is given the symbol $\mathbf{i_p}$.

Sir David Brewster discovered in 1815 that the polarising angle of a medium is fixed by its refractive index alone, through a relation now called **Brewster's law**. Its derivation is short and instructive. Take a ray of unpolarised light incident at the polarising angle i_p on the boundary between a medium of refractive index unity (air) and one of refractive index n. The reflected ray leaves at angle i_p on the same side of the normal, and the refracted ray bends into the denser medium at some angle r. Brewster's observation, which is the experimental input to the derivation, is that at the polarising angle the reflected ray and the refracted ray are perpendicular to each other. In symbols,

$$i_p + r = 90°$$

Snell's law for the boundary reads $\sin i_p = n \sin r$. Substitute $r = 90° - i_p$, so that $\sin r = \sin(90° - i_p) = \cos i_p$, and the equation becomes $\sin i_p = n \cos i_p$. Divide both sides by $\cos i_p$:

$$\mathbf{\tan i_p = n}$$

Brewster's law in its compact form. In words:

The tangent of the polarising angle of a transparent dielectric medium is equal to the refractive index of that medium relative to the medium of incidence.

Three numerical examples drive home what the law actually says. For water, $n = 1.33$; the polarising angle is $\tan^{-1}(1.33) \approx 53.1°$ measured from the normal, which is 36.9° measured from the water surface. For crown glass, $n = 1.52$; the polarising angle is $\tan^{-1}(1.52) \approx 56.7°$. For diamond, $n = 2.42$; the polarising angle is $\tan^{-1}(2.42) \approx 67.5°$. The angle creeps upward with refractive index but stays well clear of grazing incidence even for the densest natural medium normally encountered.

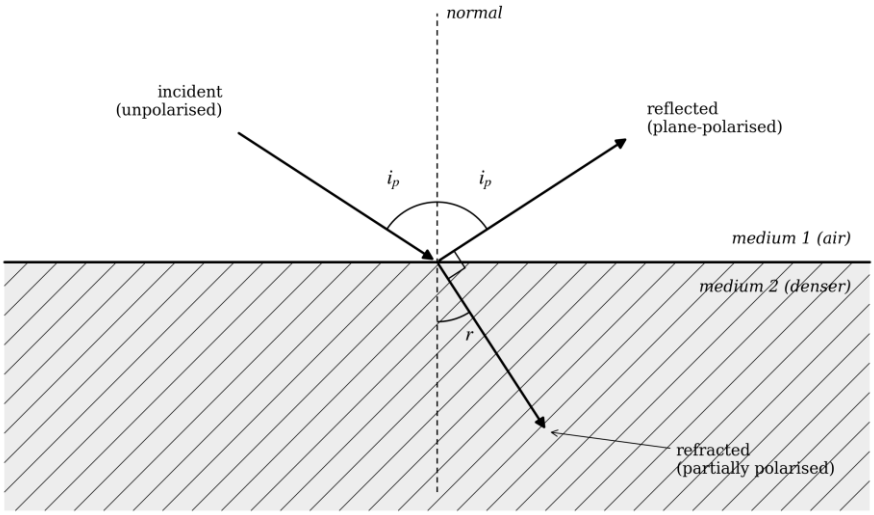


Figure: *Reflection and refraction at the polarising angle. The incident ray is unpolarised. The reflected ray, which leaves the surface at angle i_p from the normal, is completely plane-polarised with its electric vector parallel to the reflecting surface. The refracted ray, bent into the denser medium at angle r, is partially polarised. The small square at the origin marks the right angle between the reflected and refracted rays, the geometric condition that defines i_p.*

Two practical observations close the topic.

1) ***Brewster's law explains why glare from a road, the surface of a lake, or a glass shop window is partially or fully horizontal:*** it is the horizontal component of the original light that the surface preferentially reflects. Polarising sunglasses are deliberately oriented to absorb horizontal vibrations, and so dim the reflected glare without dimming everything else as much.
2) **Brewster's relation gives an easy way to measure the refractive index of an unknown sample**: place it on the bench, illuminate it at a variable angle, and find the angle at which the reflected light vanishes when viewed through an analysing polaroid whose axis is vertical. That angle is i_p; its tangent is the index.

Double refraction and the Nicol prism

Iceland spar, a clear variety of calcite ($CaCO_3$), does something unusual when light is sent through it. A single ray entering the crystal emerges as two rays. Pull the crystal across a printed page and the print is seen double. One of the two refracted rays obeys Snell's law, with a refractive index $n_o = 1.658$; the other does not, and its effective refractive index n_e depends on the direction of travel inside the crystal, ranging from 1.486 along one crystal axis to 1.658 along another. The first is called the **ordinary ray** (the o-ray); the second is called the **extraordinary ray** (the e-ray). The phenomenon as a whole is **double refraction** or **birefringence**, and it is found in any crystal whose lattice has different mechanical stiffness along different axes.

Two further facts make double refraction a useful tool.

First, *the o-ray and the e-ray are each individually plane-polarised, and their planes of vibration are at right angles to each other.* The crystal does not destroy the polarisation of the original beam; it sorts it into two beams of perpendicular polarisations.

Second, *the two beams travel in different directions inside the crystal because of their different indices, and so emerge displaced from each other.* If the crystal is thick enough and the user is patient enough, the two beams can be separated by an aperture and one of them used as a source of plane-polarised light.

William Nicol of Edinburgh, in 1828, replaced this awkward aperture trick with a piece of optical engineering of striking economy. He took a parallelepiped of Iceland spar, cut it diagonally along a plane making sixty-eight degrees with the long faces, polished the two pieces flat, and cemented them back together with Canada balsam, a clear resin of refractive index about 1.55. Unpolarised light entering one end of the cemented prism splits into the o-ray and the e-ray as before. The two rays then hit the slanting cement layer at an angle. The o-ray, with its high refractive index of 1.658 inside the calcite, is travelling from a denser medium into a less dense one (the balsam at 1.55); its angle of incidence at the cement is greater than the critical angle for that pair of indices, and it suffers total internal reflection. The e-ray, with its lower effective index of 1.486, is travelling from a less dense medium into a denser one and is transmitted through the cement into the second piece of calcite, eventually emerging at the far end. The o-ray, after total internal reflection, strikes the long side face of the prism, which Nicol blackened with a layer of paint so that the o-ray would be absorbed and not return to confuse the user.

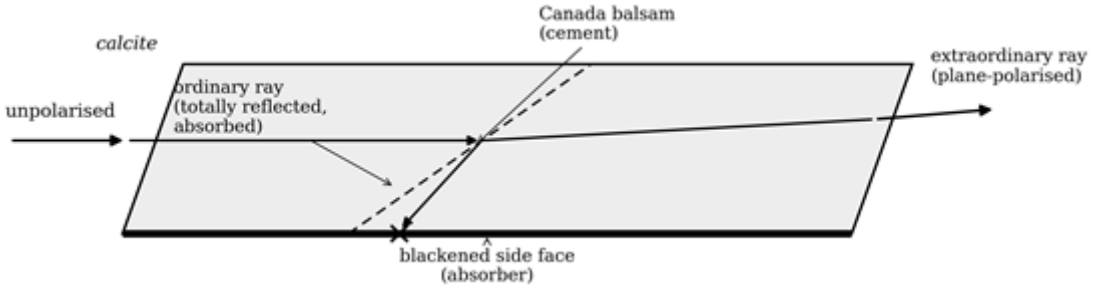


Figure: ***The Nicol prism.*** *A parallelepiped of calcite is cut diagonally and the two halves cemented back together with Canada balsam (dashed line). Unpolarised light entering the left face is split inside the crystal into an ordinary ray and an extraordinary ray, each plane-polarised. The o-ray is totally internally reflected at the cement and absorbed at the blackened side face; the e-ray passes through the cement and emerges at the right face, plane-polarised.*

What emerges from a Nicol prism is therefore a single beam of plane-polarised light, with the unwanted polarisation absorbed by the blackened side face rather than allowed to escape. Two Nicol prisms in series form a polariser-analyser pair as good as any made from Polaroid film, and historically far more accurate. The high cost of optical-quality calcite has now reduced the prism to a museum exhibit in most countries, but its design is examined because it lets the candidate apply the laws of geometrical optics (Snell's law, total internal reflection) to a real device of unmistakable elegance.

Why the sky is blue

The fourth mechanism by which unpolarised light becomes plane-polarised takes place not in a manufactured device but in the air overhead. Sunlight passing through the atmosphere is scattered by the nitrogen and oxygen molecules of the air. The wavelength of visible light is hundreds of nanometres; the size of one of these molecules is less than a nanometre. Scattering by particles much smaller than a wavelength is described by Lord Rayleigh's theory, which gives two predictions that bear directly on this section.

Rayleigh's first prediction is the **inverse-fourth-power scattering law**. Stated in words: *the intensity of light scattered sideways by a particle whose size is much smaller than the wavelength of the light is inversely proportional to the fourth power of the wavelength.* In symbols, with I_{scat} the scattered intensity and λ the wavelength,

$$\mathbf{I_{scat} \propto \frac{1}{\lambda^4}}$$

Blue light has roughly half the wavelength of red light; the factor of two in the wavelength becomes a factor of sixteen in the scattered intensity, and the sky away from the Sun appears blue because it scatters short wavelengths much more efficiently than long ones. The same argument run backwards explains the redness of the Sun at sunrise and sunset: the long path through the atmosphere strips the blue wavelengths out of the direct beam by scattering them sideways, leaving the residual direct beam dominated by the long wavelengths it could not lose.

Rayleigh's second prediction, the one that earns scattering its place in this section, is that the side-scattered light is **plane-polarised**. The mechanism is the same one that drives the polaroid film. Each air molecule is set into oscillation by the electric vector of the incoming sunlight; the oscillating charge then radiates as a tiny dipole, and a dipole radiates most strongly perpendicular to its own line of oscillation and not at all along it. Light scattered through a right angle from the Sun therefore contains only the component whose oscillation is perpendicular to the scattering direction; the parallel component cannot be radiated in that direction at all. The blue patch of sky ninety degrees from the Sun is plane-polarised, and a single polaroid filter rotated against it darkens and brightens by a factor of roughly four as the filter turns through a right angle. The student can do this with a borrowed pair of polarising sunglasses on a bright afternoon.

Optical activity

One last property of plane-polarised light is worth introducing before the section closes. A small number of substances have the curious power of rotating the plane of vibration of plane-polarised light that passes through them. Quartz crystals do it; solutions of cane sugar, tartaric acid, and a number of other organic molecules do it; many crystalline minerals do it; the antibiotic penicillin does it. Substances that have this property are called **optically active**, and the rotation is described as **dextrorotatory** (right-handed, as viewed by an observer looking back into the oncoming beam) or **laevorotatory** (left-handed) according to the sense of the rotation.

The total angle of rotation produced by a given sample is described by **Biot's law of optical activity**, first stated by Jean-Baptiste Biot in 1815. In words:

When plane-polarised light passes through a solution of an optically active substance, the angle through which the plane of vibration is rotated is directly proportional to the path length of the light through the solution and to the concentration of the solution.

Symbolically, for a solution of **concentration c** (in *grams per cubic centimetre*) and a light path of **length L** (in *decimetres*), the rotation α in degrees is:

$$\alpha = \alpha_0 L c,$$

where α_0 is the **specific rotation** of the substance, a characteristic constant determined by the **nature of the substance**. For cane sugar in water at twenty degrees Celsius and at the yellow sodium wavelength 589nm, $\alpha_0 = +66.5^\circ \text{dm}^{-1}\text{g}^{-1}\text{cm}^3$. An instrument that exploits the linearity of this relationship to determine the concentration of a sugar solution by measuring its optical rotation is called a **saccharimeter**, and is used in the sugar industry to assay the strength of cane juice and the purity of refined sugar. The student should be able to recognise that optical activity is the rotation of the plane of vibration, and to apply the linear law above to a numerical problem.

Three short problems now turn these ideas into numbers.

BINDER Example 22

Plane-polarised light of intensity 12.0Wm^{-2} falls on an analysing polaroid whose transmission axis makes an angle of 40° with the plane of vibration of the light. Calculate the intensity of the light that emerges from the analyser.

Solution

Let $I_0 = 12.0\text{Wm}^{-2}$ be the incident intensity and $\theta = 40^\circ$ the angle between the analyser axis and the plane of vibration of the incoming light. Malus's law gives the transmitted intensity directly,

$$I = I_0 \cos^2\theta.$$

$$I = 12.0\text{Wm}^{-2} \times \cos^2 40^\circ == 7.04\text{Wm}^{-2}$$

The analyser transmits 7.04Wm^{-2} of light, which is fifty-nine per cent of the incident intensity.

Making Sense of the Answer: *The number is sensible: at zero degrees the answer would be the full incident value; at ninety degrees it would be zero; the fortieth-degree value sits closer to the upper end than to the middle of that interval because forty degrees is closer to zero than it is to a right angle. A quick sanity check is that* $\cos^2 45^\circ$ *would give exactly half the intensity,* 6.00Wm^{-2}*, which it does to three significant figures.*

Think Like a Physicist: *Malus's law is the law of projection. The cosine factor in it is the projection of the incoming electric vector onto the analyser's axis; the square is there because intensity is amplitude squared. Whenever a polariser problem hands you an angle, ask yourself whether that angle is the one between the two transmission axes or between the incoming plane of vibration and one of the axes. Both ways of stating the geometry give the same answer, but only one of them is the angle that goes straight into the formula.*

REAL Example 23

Unpolarised light of intensity I_0 passes in succession through three identical polaroids whose transmission axes are oriented at 0°, 30°, and 60° to the vertical. Calculate (a) the intensity of the light that emerges from the third polaroid, and (b) the intensity that would emerge if the middle polaroid were removed. Compare the two.

Solution

Treat the journey one polaroid at a time, applying the unpolarised-light rule at the first polaroid and Malus's law at every subsequent polaroid. The angle in Malus's law is the angle between the transmission axes of consecutive polaroids, not the absolute orientation of either.

(a) The first polaroid converts unpolarised light into plane-polarised light at half the original intensity,

$$I_1 = \frac{1}{2}I_0$$

The second polaroid is at 30° to the first, so by Malus's law:

$$I_2 = I_1\cos^2(30°) = \frac{1}{2}I_0 \times \frac{3}{4} = \frac{3}{8}I_0$$

The third polaroid is at $60° - 30° = 30°$ to the second, so by Malus's law again:

$$I_3 = I_2\cos^2(30°) = \frac{3}{8}I_0 \times \frac{3}{4} = \frac{9}{32}I_0 = 0.28125I_0$$

The intensity of the light that emerges from the third polaroid is $0.28125I_0$.

(b) With the middle polaroid removed, the third polaroid in the original line-up is now the second in sequence and the angle in Malus's law becomes the full 60°,

$$I_3' = \frac{1}{2}I_0 \times \cos^2(60°) = \frac{1}{2}I_0 \times \frac{1}{4} = \frac{1}{8}I_0 = 0.125I_0$$

The intensity that would emerge if the middle polaroid were removed is $0.125I_0$.

Comparison: Adding the middle polaroid more than doubles the final intensity, from $0.125I_0$ up to $0.281I_0$.

Making Sense of the Answer: *Inserting a polaroid in the path of a polarised beam adds more light to the beam beyond it. The middle polaroid was not blocking the light; it was rotating its plane of vibration into one that the third polaroid could deal with more efficiently. Two successive thirty-degree projections lose less amplitude than one single sixty-degree projection, because $\cos 30° \times \cos 30°$ is three-quarters while $\cos 60°$ is one-half.*

Think Like a Physicist: *A polariser is a decomposer of the wave, not a curtain that lets some photons through and stops the others. Each polaroid in a series projects the incoming amplitude onto its own axis and discards the perpendicular part. Whenever a problem asks for the intensity through a chain of polaroids, walk through them one at a time, multiplying by $\cos^2\theta$ at each step, where* ***θ is the angle between the current polaroid's axis and the axis of the polaroid immediately before it.***

HOT Example 24

Light falls from air onto the calm surface of a freshwater lake whose refractive index is 1.33. Calculate (a) the polarising angle for this surface, measured both from the normal and from the surface itself, (b) the angle of refraction of the light that enters the water at this incidence, and (c) verify by direct calculation that the reflected and refracted rays are perpendicular to each other. (d) State the orientation of the electric vector in the reflected ray.

Solution

Apply Brewster's law to the air-water boundary, then Snell's law to find the refraction angle, then check the geometry.

(a) Brewster's law gives the polarising angle directly,

$$\tan i_p = n = 1.33, \quad i_p = \tan^{-1}(1.33) = 53.06° \text{(from normal)}$$

Measured from the surface itself, the polarising angle is the complement, $90° - 53.06° = 36.94°$.

(b) Snell's law for the same ray, with the air-side angle equal to i_p, reads:

$$\sin i_p = n\sin r, \quad \sin r = \frac{\sin 53.06°}{1.33} = 0.6007,$$

$$r = \sin^{-1}(0.6007) = 36.94°$$

(c) The angle between the reflected ray and the refracted ray is:

$$i_p + r = 53.06° + 36.94° = 90.00°,$$

confirming Brewster's geometric condition: at the polarising angle the reflected and refracted rays are mutually perpendicular.

(d) At Brewster incidence the reflected ray is completely plane-polarised, with its electric vector perpendicular to the plane of incidence and therefore parallel to the lake surface. The component of the original light whose electric vector lay in the plane of incidence has been fully transmitted into the water as part of the refracted ray; only the component perpendicular to that plane survives in the reflected beam.

Making Sense of the Answer: *The polarising angle for water,* $53.06°$*, is one of the three numerical anchors quoted earlier in the section and reappears here as the experimentally testable consequence of* $n = 1.33$*. The perpendicularity of reflected and refracted rays is the geometric heart of Brewster's law; it is what makes the law work, not a separate fact. The horizontal orientation of the reflected electric vector is the reason polaroid sunglasses with their absorption axis horizontal can dim the glare off the water: the very component the sunglasses are built to remove is the one that Brewster's law has just selected for the reflection.*

Think Like a Physicist: *Whenever a problem gives you a refractive index and asks for the polarising angle, write* $\tan i_p = n$ *and rearrange. The complement of* i_p *is automatically the refraction angle* r*, because Brewster's condition forces* $i_p + r = 90°$*. Two parts of the problem can therefore be answered without Snell's law at all; Snell's law is then a check on the arithmetic, not an independent step.*

Three sections of the chapter have now studied what the wave does. It interferes when two of itself meet; it diffracts when one of itself passes an obstacle; it polarises when the direction of its transverse vibration is filtered. The pages that follow take the same physics out of the laboratory and into the world. The thin films that colour a soap bubble and the anti-reflection coating on a camera lens are interference. The compact disc that throws a spectrum across the table when held under a lamp is diffraction. The screen of a calculator and the lens of a pair of sunglasses are polarisation. Each is, in its own way, an answer to a question of engineering that this chapter has now equipped the candidate to ask.

APPLICATIONS OF PHYSICAL OPTICS

So far, the chapter has built the theory of physical optics through definitions, derivations, and worked examples. This section now turns from the page to the world beyond it. It asks what these results actually do once they leave the classroom and enter everyday life. The purpose of the final section is therefore not to introduce new physics, but to reveal how remarkably far the physics already in hand can reach.

1. Holography

A hologram is a photographic record of the interference pattern between two coherent beams of light. The first beam is the object beam, scattered from the surface of the object that is being recorded; the second is the reference beam, taken straight from the source. Both come from a single laser, so their phase relationship is fixed at every point on the photographic plate. The plate records, at every pixel, the fringe contrast that the two beams produce locally. The result is not an image of the object but an interference pattern that encodes every wavefront the object would have sent toward the plate.

To replay the hologram, the developed plate is illuminated by a reproduction of the original reference beam. The plate now acts as a diffraction grating whose pitch varies from point to point exactly as the recorded fringes vary; the diffracted light reproduces the object's wavefront, and an observer on the far side sees a three-dimensional image floating in the position the original object occupied. The technique exploits **coherence** (*to record clean fringes*), **interference** (*to make the recording*), and **diffraction** (*to read it out*), and so brings together three sections of this chapter into one device. The same principle, scaled down to a few millimetres, makes the security holograms printed on bank notes and credit cards: they cannot be photocopied, because a photocopier captures an image rather than the wavefront that produced the image.

2. Anti-reflection coatings

A bare glass surface reflects about four per cent of the light incident on it. A modern camera zoom lens has fifteen or more such surfaces; without intervention the light reaching the sensor would be reduced to less than half of what entered the front element, with much of the lost light bouncing around inside the lens barrel and producing flare. The intervention is a thin coating of a transparent material, magnesium fluoride being the standard, whose refractive index 1.38 sits between the index of air and the index of crown glass at 1.52. A film of this material of thickness one quarter of the wavelength inside the film,

$$t = \frac{\lambda}{4n},$$

makes the two reflections, one from the top surface and one from the bottom, arrive in antiphase at every point above the lens. They cancel by destructive interference; the suppressed reflection is added to the transmitted beam, and the lens transmits ninety-five per cent of the light it receives instead of seventy. The same coating is used on spectacle lenses, on the cover glass of solar panels, and on the front surface of every modern smartphone screen, with the design wavelength chosen at the centre of the visible band so that the residual reflection is faint and purplish rather than at any particular spectral peak.

3. Glare-reducing polaroid sunglasses

Light reflected from a horizontal surface at large angles of incidence (a wet road, the bonnet of a car, the surface of a lake, a snowfield) is partially polarised, with its electric vector preferentially horizontal. At the Brewster angle the polarisation is complete; at angles either side of it the polarisation is substantial. A pair of Polaroid sunglasses whose transmission axes are oriented vertically absorbs the horizontal component and lets through only the vertical one. The reflected glare is dimmed by a factor of three or four; the rest of the scene, whose light is on average unpolarised, is dimmed only by the constant factor $1/2$ that any single polaroid imposes on unpolarised light.

The behaviour of the lenses when tilted is the test by which a wearer can confirm that the glasses really are polaroid rather than tinted. Tilt the head sideways by ninety degrees, and the absorption axis is now vertical; the previously dim glare returns, and the rest of the scene is roughly unchanged. Fishermen and motorists use the effect routinely; landscape painters and photographers use the same effect, mounted on a camera lens instead of on a pair of spectacles, to darken the blue sky against the clouds.

4. Photoelastic stress analysis

Many transparent plastics, when bent or compressed, become temporarily birefringent. Their refractive index for light polarised parallel to the local stress direction differs from that for light polarised perpendicular to it. A stressed plastic part viewed between two crossed Polaroid filters develops a pattern of bright and dark coloured fringes, in which the colour at each point is the wavelength for which the two polarisations have accumulated a relative phase difference of an odd integer number of half-wavelengths across the thickness of the part. The fringe pattern maps directly onto the local stress field.

Aircraft-window manufacturers, dam designers, and the makers of large plastic optical components used the technique to find stress concentrations in scale models of their parts long before the development of finite-element computer modelling. The technique survives today as the cheapest and most direct way to demonstrate to a student that stress in a plastic is a genuine physical quantity. A transparent plastic ruler bent gently between crossed Polaroid filters held up to the light shows the fringes immediately; the heaviest fringe concentration sits at the corners of any notches or holes in the ruler, exactly where an engineer would expect stress to concentrate.

5. Three-dimensional cinema and stereo viewers

A three-dimensional cinema projects two slightly different images of the same scene onto the screen at the same time, one for each eye. The two images are polarised in perpendicular linear planes, or, in the modern variant, in opposite circular polarisations. The viewer wears spectacles whose left and right lenses are Polaroid filters of perpendicular orientations (or circular polarisers of opposite handedness). Each lens passes only the image meant for its eye and blocks the other; the brain interprets the resulting pair of slightly different images as a single three-dimensional scene, exactly as it does with the two slightly different images delivered by the eyes of an ordinary observer.

The trick relies on the two perpendicular polarisations surviving the trip from projector to screen to filter. An ordinary white screen depolarises the light it scatters, because every microscopic facet of paint scatters in a random orientation; modern three-dimensional cinemas therefore use silvered screens, whose flat metallic surfaces preserve the polarisation of the light reflected from them. The same physics drives small hand-held stereo viewers that combine a left-eye and a right-eye image on a single transparency between two polarising filters.

6. X-ray diffraction

The grating equation continues to hold at much shorter wavelengths than the visible band, provided a much smaller slit spacing is available. The lattice planes of a crystal, with spacings of about $0.1\,\text{nm}$, are such a grating, and the wavelengths that can be diffracted by them are X-rays, of comparable order of magnitude. The grating equation, written for reflection off adjacent lattice planes a distance d apart, becomes **Bragg's law**, $2d\sin\theta = n\lambda$, where θ is the glancing angle of the X-ray to the lattice planes and n is the order of the maximum. The diffraction pattern that comes off a crystal is therefore a fingerprint of the crystal's lattice spacings, and from a sufficiently complete pattern, in three dimensions, the positions of every atom in the unit cell can be deduced.

The technique determined the structure of common salt in 1913, the structure of penicillin in 1945, and, most famously of all, the double-helix structure of DNA. Photograph 51, taken by Rosalind Franklin at King's College London in May 1952, is an X-ray diffraction pattern of a fibre of DNA; the characteristic cross of streaks at the centre of the pattern is the diffraction signature of a helix, and the spacing of the streaks gave Crick and Watson the helix's pitch. X-ray crystallography is, with some justification, the most consequential single application of diffraction physics in the history of science.

7. Optical-fibre interference filters

Modern telecommunications send data along optical fibres by encoding several independent streams on slightly different wavelengths of infrared light and combining them onto the same fibre. At the receiving end of the fibre the streams have to be separated. The standard device for the separation is a stack of dielectric thin films deposited on a glass substrate; the films are designed so that, by interference between their many internal reflections, the stack transmits a narrow band of wavelengths and reflects everything else. A different filter, with a different stack design, is placed in the path for each independent channel.

The technique is called wavelength-division multiplexing. A single fibre crossing the Indian Ocean can carry a hundred or more independent channels in this way, each carrying tens of gigabits of data per second; without the technique, every channel would require its own fibre, and the cost of a transatlantic cable would be a hundred times what it is. The thin films at each end of the cable are designed using exactly the quarter-wave logic that earlier in this section produced the anti-reflection coating, scaled and stacked to make passbands rather than zero-reflection points.

8. Liquid-crystal displays

The dominant electronic display technology of the modern world is built around a sandwich of two crossed Polaroid films with a liquid-crystal layer between them. The liquid-crystal molecules, when no voltage is applied across the cell, twist the plane of polarisation by ninety degrees as the light passes through them. This rotation matches the angle between the two Polaroid axes, and the cell looks bright. When a voltage is applied across the cell, the molecules untwist; the plane of polarisation passes through unchanged, the second Polaroid blocks it, and the cell looks dark.

Every pixel of every laptop screen, smartphone display, calculator, and digital watch is one such cell with its voltage independently controllable. A colour display adds three coloured filters, red, green, and blue, behind the pixel array; varying the voltages independently in three sub-pixels per pixel produces a full-colour image. The whole technology is, in effect, a programmable curtain made of Polaroid, in which the rotating part of the curtain is a fluid whose molecules respond to electric fields. Every other physical-optics phenomenon studied in this chapter is also present somewhere in the device, but the polarisation rotation is the piece that does the central work.

Now that each section has been understood individually, it is time to see the ideas interact. The following miscellaneous worked examples combine concepts from across the chapter into single, integrated problems.

MISCELLANEOUS WORKED EXAMPLES ON PHYSICAL OPTICS

Example 25

(a) Explain in physical terms why a true double-slit interference pattern is obtained when the two slits are illuminated by light from a single source through a narrow slit, but not when the same two slits are illuminated by two physically distinct lamps of the same nominal colour.

(b) Light of coherence length 0.6mm illuminates a thin oil film of refractive index 1.50. Determine the maximum film thickness above which interference colours can no longer be observed in reflected light.

Solution

(a) A stable interference pattern requires that the two slit-illuminating sources stay in a fixed phase relationship over a time long compared with the eye's or detector's integration time. Light from a single primary source split into two paths by a pair of slits originates from the same emission events; the phase relationship between the two paths is therefore the constant zero (apart from path-length differences), and the pattern is stable. Two physically separate lamps emit independent sequences of wavetrains, each lasting only a few nanoseconds. The phase relationship between the two lamps shifts randomly on this nanosecond timescale, while the detector integrates over much longer times. The resulting pattern is averaged to a uniform illumination, as the lack of coherence between the two lamps washes the fringes out.

(b) For interference to be visible from a thin film in reflected light, the optical path difference between the front-reflected and back-reflected beams must lie within the coherence length of the source. The path difference for normal incidence is 2nt, where n is the refractive index of the film and t its thickness. The maximum thickness for visible fringes is therefore set by:

$$2\text{nt}_{\text{max}} = \text{L}_\text{c}$$

$$\text{t}_{\text{max}} = \frac{\text{L}_\text{c}}{2\text{n}} = \frac{0.6\text{mm}}{2(1.50)} = 0.2\text{mm}$$

Above 0.2mm the front- and back-reflected beams are too far separated in time for interference colours to be observed.

Example 26

(a) Account in physical terms for the experimental observation that the central fringe of a Young's double-slit pattern remains bright when the entire apparatus is immersed in water, while every other fringe moves closer to the centre.

(b) In a Young's double-slit experiment, light of wavelength 600nm illuminates two slits 0.40mm apart, and fringes are observed on a screen 1.50m away. Find the fringe width on the screen, and the distance between the third bright fringe and the third dark fringe on the same side of the central fringe.

Solution

(a) The central fringe corresponds to zero path difference between light reaching the screen from the two slits. Submerging the apparatus in water does not alter the geometry of the two paths; both paths run through the same medium and remain equal in length, so the central fringe stays at the same position on the screen and remains bright.

Higher-order fringes occur at path differences equal to integer multiples of the wavelength. The wavelength of light inside water is smaller than the in-air value by a factor equal to the refractive index of water. The fringe spacing $\beta = \lambda\text{D}/\text{d}$ shrinks by the same factor. All non-central fringes therefore move closer to the centre, with the central fringe anchored as the only fixed reference.

(b) Fringe width:

$$\beta = \frac{\lambda\text{D}}{\text{d}} = \frac{(600\times10^{-9}\text{m})(1.50\text{m})}{0.40\times10^{-3}\text{m}} = 2.25\times10^{-3}\text{m} = 2.25\text{mm}$$

Position of the third bright fringe: $\text{y}_{3,\text{br}} = 3\beta = 6.75\text{mm}$.

Position of the third dark fringe (third dark counted from the central bright, at half-integer fringe positions): $y_{3,dk} = (3 - 1/2)\beta = 2.5\beta = 5.625\text{mm}$.
Distance between them:

$$\Delta y = y_{3,br} - y_{3,dk} = 6.75\text{mm} - 5.625\text{mm} = 1.125\text{mm} = \frac{\beta}{2}$$

Example 27

(a) Explain in physical terms why, when a Young's double-slit experiment is performed with white light, the central fringe is white but higher-order fringes show coloured fringing.
(b) Two wavelengths, 650nm and 520nm, simultaneously illuminate a Young's double-slit whose slit separation is 0.50mm and screen distance is 1.20m. Find the distance from the central fringe to the first position at which a bright fringe of each wavelength coincides.

Solution

(a) The central fringe is at zero path difference. At zero path difference every wavelength interferes constructively, so all the wavelengths in white light arrive at the central fringe in phase and add. The combined fringe is therefore bright and white.

Away from the centre, the constructive condition becomes path difference equal to an integer number of wavelengths. The wavelength entering this condition is different for each colour, so the first bright fringe of red light is farther from the centre than that of violet light. Consecutive first-order fringes from different wavelengths therefore fan out into a band of colour, with violet inside and red outside the central white fringe. Higher orders overlap and lose distinct colour identity.

(b) Coincidence of bright fringes requires $n_1\lambda_1 = n_2\lambda_2$ for integers n_1 and n_2.

$$\frac{n_1}{n_2} = \frac{\lambda_2}{\lambda_1} = \frac{520}{650} = \frac{4}{5}$$

The smallest such integers are $n_1 = 4$, $n_2 = 5$ (and $4 \times 650\text{nm} = 5 \times 520\text{nm} = 2600\text{nm}$).
Position on the screen:

$$y = \frac{n_1\lambda_1 D}{d} = \frac{4(650 \times 10^{-9}\text{m})(1.20\text{m})}{0.50 \times 10^{-3}\text{m}} = 6.24 \times 10^{-3}\text{m} = 6.24\text{mm}.$$

Example 28

(a) Distinguish between the appearance of the pattern observed on the screen in a Young's experiment when one of the two slits is covered with an opaque card and when both slits are open.
(b) In a Young's double-slit experiment, the slit separation is 0.30mm and the screen is at a distance of 1.20m from the slits. The wavelength of the illuminating laser is changed from 632.8nm (red) to 532nm (green). Find the change in fringe width.

Solution

(a) When one slit is covered, only the uncovered slit acts as a source. The screen shows a single-slit diffraction pattern: a wide central bright maximum, flanked by progressively weaker secondary maxima separated by dark minima. No interference fringes are present, because there is no second source whose light can superpose with that of the first.

When both slits are open, light from each slit reaches every point on the screen and superposes. The screen now shows the Young's interference fringes of equal spacing $\beta = \lambda D/d$, modulated by the single-slit diffraction envelope from each slit. The interference fringes are evenly spaced, the diffraction envelope is broad and bell-shaped, and the visible pattern is a series of evenly spaced bright bars whose intensity decreases away from the centre according to the diffraction envelope.

(b) Fringe width depends on wavelength as $\beta = \lambda D/d$, so the change in fringe width when the wavelength changes from λ_1 to λ_2 is:

$$\Delta\beta = \frac{D(\lambda_1 - \lambda_2)}{d} = \frac{(1.20\text{m})(632.8 \times 10^{-9}\text{m} - 532 \times 10^{-9}\text{m})}{0.30 \times 10^{-3}\text{m}} = 4.03 \times 10^{-4}\text{m} \approx 0.40\text{mm}.$$

The fringe width decreases by approximately 0.40mm when the laser changes from red to green.

Example 29

(a) Explain in physical terms why placing a thin transparent sheet over one of the two slits in a Young's double-slit experiment shifts the entire fringe pattern toward the side covered by the sheet.
(b) In a Young's double-slit experiment with slit separation 0.40mm and screen distance 1.50m, illuminated by light of wavelength 550nm, a thin mica sheet of refractive index 1.58 is placed over one of the slits. The central fringe shifts on the screen by 0.20cm. Determine the thickness of the mica sheet.

Solution

(a) Light passing through a thin transparent sheet of refractive index μ and thickness t covers an optical path length μt inside the sheet, whereas in the absence of the sheet the geometric path length over the same distance would be t. The light traversing the covered slit therefore accumulates an extra optical path of $(\mu - 1)t$ compared with the light through the uncovered slit.

To re-establish zero net path difference between the two beams (which defines the central bright fringe), the geometric path on the covered side must shorten by the same amount. The central fringe therefore moves toward the covered side, and the entire pattern shifts with it.

(b) Fringe shift formula:

$$\Delta y = \frac{D(\mu - 1)t}{d}$$

Solving for t:

$$t = \frac{\Delta y \times d}{D(\mu - 1)} = \frac{(0.20 \times 10^{-2}\text{m})(0.40 \times 10^{-3}\text{m})}{(1.50\text{m})(1.58 - 1)} = 9.2 \times 10^{-7}\text{m} = 0.92\ \mu\text{m}$$

Example 30

(a) Account in physical terms for the changing pattern of colours that moves across the surface of a soap film stretched in a wire ring as the film drains under gravity.

(b) A soap film of refractive index 1.33 is illuminated normally by white light. Find the smallest film thickness that produces a bright reflected fringe at a wavelength of 500nm.

Solution

(a) A soap film stretched in a wire ring drains under gravity, becoming thicker at the bottom and thinner at the top. The thickness varies smoothly and continuously with height. Light incident on the film reflects from both the front surface and the back surface; the back-reflected beam travels an extra optical path of $2nt$ inside the film (with an additional half-wavelength shift due to the π phase change on reflection from the front surface). The condition for bright reflection therefore selects, at each thickness, a particular wavelength satisfying $2nt = (m - 1/2)\lambda$. Because t varies along the height of the film, the wavelength satisfying the bright condition also varies smoothly, producing the bands of colour that flow visibly down the film as it drains.

(b) For bright reflection from a thin film of refractive index n and thickness t at wavelength λ, with the half-wavelength shift on reflection accounted for:

$$2nt = (m + 1/2)\lambda, \qquad m = 0,1,2, \ldots$$

Smallest thickness corresponds to $m = 0$:

$$t_{min} = \frac{\lambda}{4n} = \frac{500 \times 10^{-9}\text{m}}{4(1.33)} = 9.40 \times 10^{-8}\text{m} = 94.0\text{nm}$$

Example 31

(a) A student wishes to coat a glass lens of refractive index 1.50 with magnesium fluoride of refractive index 1.38 so that the intensity of light reflected at a wavelength of 550nm is reduced. Account in physical terms for why magnesium fluoride works as an anti-reflection coating on glass.

(b) A piece of thin copper wire is laid between two flat glass plates at one end to form a small air wedge. When the wedge is illuminated normally with sodium light of wavelength 589nm, a total of 102 dark fringes are counted across the entire length of the wedge, from the closed end to the open end where the wire sits. Calculate the diameter of the wire.

Solution

(a) Light incident on the coated lens reflects partially at the air–coating interface and partially at the coating–glass interface. Both reflections are from a lower-index medium into a higher-index medium, so both undergo a π phase change on reflection. The two π shifts cancel, leaving the relative phase of the two reflected beams to be set entirely by the optical path difference $2nt$. Choosing the coating thickness so that $2nt$ equals an odd half-wavelength produces destructive interference between the two reflections, and the net reflected intensity is minimised. The refractive index of magnesium fluoride (1.38) lies between that of air (1.00) and that of glass (1.50), so the two interface reflections have comparable amplitude, which is a requirement for nearly complete cancellation.

(b) The wire's diameter equals the thickness of the air wedge at the open end. The total number N of dark fringes spanning the wedge from t = 0 at the closed end to t = h at the open end is given by:

$$N = 2h / \lambda$$

Solving for h:

$$h = N\lambda / 2$$

Substituting N = 102 and λ = 589nm:

$$h = (102 \times 589 \times 10^{-9}\ \text{m}) / 2 = 3.00 \times 10^{-5}\ \text{m} = 30.0\mu\text{m} = 0.030\text{mm}$$

The diameter of the copper wire is therefore 0.030mm, or 30μm.

Example 32

(a) Account for the experimental fact that the centre of the Newton's-rings pattern observed in reflected light is dark, not bright.

(b) In a Newton's-rings experiment, the diameter of the eighth dark ring (measured from the centre) is 4.80mm and the radius of curvature of the lens is 1.20m. Calculate the wavelength of the light used.

Solution

(a) At the centre of the Newton's-rings arrangement the lens makes contact with the flat glass plate, and the air gap between them is essentially zero. Light reflecting from the lower surface of the lens (going from glass into air) undergoes no phase change, because it is reflecting from a less-dense medium. Light reflecting from the upper surface of the glass plate (going from air into glass) undergoes a π phase change, because it is reflecting from a denser medium. Even though the geometric path difference between the two reflections is zero at the centre, the phase difference is π. The two beams therefore interfere destructively, and the centre is dark.

(b) For dark rings in Newton's rings (reflected light), the diameter of the nth dark ring satisfies

$$D_n^2 = 4nR\lambda,$$

where R is the radius of curvature of the lens. Solving for λ:

$$\lambda = \frac{D_n^2}{4nR} = \frac{(4.80\times10^{-3}\text{m})^2}{4(8)(1.20\text{m})} = 6.00\times10^{-7}\text{m} = 600\text{nm}$$

Example 33

(a) Account in physical terms for the fact that, in a Newton's-rings experiment, a plot of the squared diameter of successive dark rings against the ring number gives a straight line through the origin.

(b) In a Newton's-rings experiment, the squared diameter of the twelfth dark ring exceeds that of the sixth dark ring by 12.0mm^2. The radius of curvature of the lens is 0.80m. Calculate the wavelength of the light used.

Solution

(a) The condition for the nth dark ring in reflected light is that the air-gap thickness t at that radius gives a path difference satisfying $2t = n\lambda$. The geometry of a spherical lens resting on a flat plate gives, for small t, $t \approx r^2/(2R)$, where r is the radius of the ring and R is the radius of curvature of the lens. Substituting:

$$\frac{r^2}{R} = n\lambda \qquad \Rightarrow \qquad D_n^2 = 4r^2 = 4nR\lambda$$

This is a linear relation between D_n^2 and n, with slope $4R\lambda$ and zero intercept. A plot of D_n^2 versus n is therefore a straight line through the origin (and either of R and λ may be determined graphically).

(b) From $D_n^2 = 4nR\lambda$, the difference of squared diameters of two rings at orders $n_2 > n_1$ is:

$$D_{n_2}^2 - D_{n_1}^2 = 4R\lambda(n_2 - n_1).$$

For $n_2 = 12$, $n_1 = 6$:

$$\lambda = \frac{D_{12}^2 - D_6^2}{4R(12-6)} = \frac{12.0\times10^{-6}\text{m}^2}{24(0.80\text{m})} = 6.25\times10^{-7}\text{m} = 625\text{nm}$$

Example 34

(a) Distinguish between the diffraction pattern produced by a single slit and the interference pattern produced by a Young's double-slit, with reference to the spacing and intensity of consecutive bright fringes.

(b) A monochromatic beam of wavelength 600nm passes through a single slit of width 0.10mm and forms a pattern on a screen 2.0m away. Find the width of the central maximum, and the angular position of the second-order minimum.

Solution

(a) The single-slit diffraction pattern consists of a wide central bright maximum flanked by progressively narrower secondary maxima separated by dark minima. The central maximum is roughly twice the width of the secondary maxima, and the intensity of the secondary maxima falls off rapidly (the first secondary maximum is only about 4.5% of the central maximum).

The Young's double-slit interference pattern consists of equally spaced bright fringes of nearly equal intensity, modulated by the single-slit diffraction envelope of each individual slit. Successive bright fringes are evenly spaced (in contrast to the single-slit pattern's narrowing secondary maxima), and the intensity of consecutive bright fringes is nearly equal across the central region of the envelope.

(b) The central maximum extends between the first dark minima on either side. The first minimum lies at $a\sin\theta = \lambda$, so $\sin\theta \approx \lambda/a$ and the half-width on the screen is $y_1 \approx D\lambda/a$. The full width is:

$$W = 2y_1 = \frac{2D\lambda}{a} = \frac{2(2.0\text{m})(600\times10^{-9}\text{m})}{0.10\times10^{-3}\text{m}} = 2.4\times10^{-2}\text{m} = 24\text{mm}$$

Angular position of the second-order minimum:

$$\sin\theta_2 = \frac{2\lambda}{a} = \frac{2(600\times10^{-9})}{0.10\times10^{-3}} = 1.2\times10^{-2} \qquad \Rightarrow \qquad \theta_2 = \sin^{-1}(1.2\times10^{-2}) = 0.69°.$$

Example 35

(a) Account for the experimental observation that, no matter how much the line density of a diffraction grating is increased, only a finite number of principal maxima can ever be observed for a given wavelength of light.

(b) A diffraction grating has 500 lines per millimetre. Light of wavelength 580nm is incident normally on the grating. Find all the diffraction orders that can be observed, and the diffraction angle of the highest visible order.

Solution

(a) The grating equation $d\sin\theta = n\lambda$ requires $\sin\theta = n\lambda/d$ for the nth principal maximum, where d is the grating spacing. The function $\sin\theta$ cannot exceed unity, so the order number is bounded by

$$n \leq \frac{d}{\lambda}.$$

A grating with a higher line density has a smaller d, and a smaller d means a smaller value of d/λ and therefore a smaller maximum order. Increasing the line density indefinitely makes d approach λ from above; below that, no orders beyond the central maximum exist at all. For any given wavelength, the number of observable principal maxima is therefore finite, and the upper bound is set by the ratio d/λ.

(b) Grating spacing:

$$d = \frac{1}{500 \times 10^3 m^{-1}} = 2.0 \times 10^{-6} m.$$

Maximum order:

$$n_{max} = \left\lfloor \frac{d}{\lambda} \right\rfloor = \left\lfloor \frac{2.0 \times 10^{-6}}{5.80 \times 10^{-7}} \right\rfloor = \lfloor 3.45 \rfloor = 3.$$

Observable orders: $n = 0, \pm 1, \pm 2, \pm 3$. The highest visible order is the third.

Diffraction angle of the third order:

$$\sin\theta_3 = \frac{3\lambda}{d} = \frac{3(5.80 \times 10^{-7})}{2.0 \times 10^{-6}} = 0.870 \qquad \Rightarrow \qquad \theta_3 = \sin^{-1}(0.870) = 60.5°.$$

Example 36

(a) Explain in physical terms why a grating made of equally spaced opaque wires separated by transparent gaps produces a diffraction pattern whose principal maxima lie at the same angles as those produced by a grating made of equally spaced transparent rulings on glass.

(b) Infrared light of wavelength $1.0 \times 10^{-5} m$ is incident normally on a grating made of fine wires of diameter $1.0 \times 10^{-5} m$ with transparent gaps of width $1.5 \times 10^{-5} m$ between consecutive wires. Find the angle of the first-order maximum.

Solution

(a) The grating equation $d\sin\theta = n\lambda$ relates the diffraction angles of the principal maxima to the period d of the grating, defined as the distance from one feature to the next identical feature. Whether each period consists of one opaque wire and one transparent gap, or one transparent ruling and one opaque gap, the period is the same: the centre-to-centre spacing of the wires equals the centre-to-centre spacing of the corresponding transparent rulings on the complementary grating. The principal maxima therefore occur at the same angles in the two cases. The two gratings differ in the relative intensities of the orders (because the diffraction envelope from a single opaque wire is not the same as from a single transparent ruling), but not in the angular positions of the principal maxima.

(b) Grating period:

$$d = d_{wire} + d_{gap} = 1.0 \times 10^{-5} m + 1.5 \times 10^{-5} m = 2.5 \times 10^{-5} m$$

First-order angle:

$$\sin\theta_1 = \frac{\lambda}{d} = \frac{1.0 \times 10^{-5}}{2.5 \times 10^{-5}} = 0.40 \qquad \Rightarrow \qquad \theta_1 = 23.6°$$

Example 37

(a) Account in physical terms for the fact that sound waves in air cannot be polarised, while light waves can.

(b) Unpolarised light of intensity $16 W/m^2$ passes through a Polaroid sheet, then through a second Polaroid sheet whose transmission axis is at 30° to that of the first. Find the intensity of the light emerging from the second sheet.

Solution

(a) Polarisation is the restriction of a wave's vibration to a specified plane perpendicular to the direction of propagation. Sound waves in air are longitudinal: the air molecules oscillate along the direction of propagation, so there is no plane transverse to propagation in which the vibration could be restricted. Polarisation is meaningless for a longitudinal wave. Light waves are transverse: the electric (and magnetic) field oscillates in a plane perpendicular to the direction of propagation, and that plane has two independent directions of oscillation. Restricting the electric vector to a single direction in that plane (by a Polaroid filter, by Brewster reflection, by selective scattering, and so on) defines a polarised state. The contrast between sound and light is a direct consequence of their longitudinal versus transverse character.

(b) Unpolarised light passing through the first Polaroid:

$$I_1 = \frac{I_0}{2} = \frac{16W/m^2}{2} = 8W/m^2$$

The transmitted light is linearly polarised along the first sheet's axis. Through the second Polaroid (Malus's law):

$$I_2 = I_1\cos^2\theta = (8W/m^2)\cos^2(30° = 6W/m^2$$

Example 38

(a) Account in physical terms for the observation that inserting a third Polaroid sheet between two crossed Polaroids causes light to be transmitted through the combination, even though no light passes through the two crossed Polaroids alone.

(b) Unpolarised light of intensity $24W/m^2$ is incident on three Polaroid sheets stacked in series. The first and third are crossed (their transmission axes are at $90°$ to each other); the second is placed midway between them with its transmission axis at $45°$ to each of the others. Find the intensity of the light emerging from the third sheet.

Solution

(a) Two crossed Polaroids transmit no light because the second sheet is oriented to block exactly the polarisation that the first sheet transmits. When a third Polaroid is inserted between them at, say, $45°$ to the axis of the first, the situation changes. After the first Polaroid, the light is polarised along axis-1. The third (middle) Polaroid, at $45°$ to axis-1, transmits the component along its own axis (the cosine-squared factor at $45°$, i.e., one half), and re-polarises the light along axis-3. Light now reaches the last Polaroid polarised along axis-3, which is at $45°$ to the last polaroid's axis. The last Polaroid transmits another factor of one half. Net: the previously blocked light now reaches the screen because the inserted Polaroid has rotated the polarisation in steps of $45°$ rather than blocking it outright.

(b) After the first Polaroid:

$$I_1 = \frac{I_0}{2} = \frac{24W/m^2}{2} = 12W/m^2$$

After the middle Polaroid (at $45°$ to the first):

$$I_2 = I_1\cos^2(45°) = (12W/m^2)(0.5) = 6W/m^2$$

After the third Polaroid (at $90°$ to the first, i.e., $45°$ to the middle):

$$I_3 = I_2\cos^2(45°) = (6W/m^2)(0.5) = 3W/m^2$$

Example 39

(a) Account in physical terms for the experimental observation that Polaroid sunglasses with a vertical transmission axis reduce the glare of sunlight reflected from a wet road but not the glare of sunlight reflected from a polished mirror at normal incidence.

(b) Sunlight reflects from the surface of a calm lake of refractive index 1.33 at the Brewster angle. Find the angle of the reflected ray above the horizontal, the angle of refraction in the lake, and the orientation in space of the electric vector of the reflected light.

Solution

(a) Light reflected from a horizontal surface at a large angle of incidence is partially polarised in the horizontal direction (the component of the electric vector parallel to the surface is preferentially reflected). Polaroid sunglasses with a vertical transmission axis block most of this horizontally polarised glare from the wet road, sharply reducing the perceived brightness of the reflection.

Light reflected from a polished mirror at normal incidence is, however, unpolarised. At normal incidence, the two orthogonal components of the electric field are reflected with equal intensity, and the reflected light shows no preferred direction of vibration. Polaroid sunglasses transmit only one component (half the intensity) regardless of orientation, so they do not selectively suppress the mirror reflection; the sunglasses simply dim the reflection by half.

(b) Brewster angle:

$$\tan i_p = n_{water} = 1.33 \qquad \Rightarrow \qquad i_p = \tan^{-1}(1.33) = 53.06°$$

Angle of reflection above the horizontal: $90° - i_p = 36.9°$

At the Brewster angle, the refracted ray is exactly perpendicular to the reflected ray, so:

$$r = 90° - i_p = 36.9°$$

The electric vector of the reflected light is perpendicular to the plane of incidence, i.e., parallel to the surface of the lake (horizontal and along the line of the shore).

Example 40

(a) A clear plastic ruler placed between two crossed Polaroid sheets shows a delicate, multi-coloured pattern when bent slightly. Account in physical terms for this observation.

(b) X-rays of unknown wavelength are reflected by a crystal of sodium chloride whose interplanar spacing is $0.282nm$. The first-order Bragg reflection is observed at a glancing angle of $7.5°$. Find the wavelength of the X-rays.

Solution

(a) Most transparent solids are optically isotropic, so light propagates through them at a single refractive index regardless of polarisation. When such a solid is mechanically stressed, the strain field breaks the isotropy and makes the material

birefringent: the refractive index for light polarised along the stress direction is now slightly different from that for light polarised at right angles. Plane-polarised light passing through such a stressed region therefore emerges elliptically polarised, with the ellipse's shape and orientation depending on the local stress. When this elliptical light is viewed through the analyser (the second Polaroid, crossed with the first), the transmitted intensity depends on the elliptical state, and therefore on the local stress. White-light illumination produces a different colour for each combination of strain magnitude and direction, mapping the stress field across the ruler into a coloured pattern of fringes. This is the principle of photoelastic stress analysis.

(b) Bragg's law for the nth-order reflection:

$$2d\sin\theta = n\lambda$$

For $n = 1$, $d = 0.282\text{nm}$, $\theta = 7.5°$:

$$\lambda = 2d\sin\theta = 2(0.282 \times 10^{-9}\text{m})\sin 7.5° = 7.36 \times 10^{-11}\text{m}$$

The wavelength of the X-rays is approximately $7.36 \times 10^{-11}\text{m} = 0.0736\text{nm}$.

Example 41

(a) Account in physical terms for the experimental observation that a hologram of an object recorded with red laser light reproduces a sharp three-dimensional virtual image only when the hologram is illuminated by light of the same wavelength as that used to record it.

(b) A laser used to make a hologram has a wavelength of 632.8nm in air. The hologram is then submerged in water of refractive index 1.33. Find the wavelength of the laser light inside the water.

Solution

(a) A hologram is a finely detailed interference pattern recorded in a photosensitive emulsion. The pattern is the result of interference between a reference beam and the light scattered from the object during recording. The recorded fringe structure encodes both the amplitude and the phase of the object wavefronts. When the hologram is later illuminated with the original reference beam (same wavelength), the recorded fringe pattern diffracts the reference beam in the geometric pattern of a grating, reconstructing the original object wavefronts and producing a sharp three-dimensional virtual image.

Diffraction from any periodic structure follows $d\sin\theta = n\lambda$, in which the angles depend on λ. Illuminating the same hologram with a different wavelength changes all the diffraction angles, so the reconstructed wavefronts no longer match the original object geometry. The image is correspondingly distorted or blurred. Only the original wavelength produces a sharp reconstruction.

(b) The wavelength of light inside a medium of refractive index n is:

$$\lambda_{\text{water}} = \frac{\lambda_{\text{air}}}{n_{\text{water}}} = \frac{632.8\text{nm}}{1.33} = 475.8\text{nm}$$

Example 42

(a) Explain in physical terms why a beam of plane-polarised light, after passing through a column of sugar solution, has its plane of polarisation rotated.

(b) A column of sugar solution of length 10cm and concentration 0.20g/cm^3 rotates the plane of polarisation of yellow light by $13°$. Calculate the specific rotation of the sugar at this wavelength.

Solution

(a) Sugar molecules are chiral: their three-dimensional structure has a definite handedness (sugar is right-handed). A plane-polarised wave entering the solution can be decomposed into two circularly polarised components of equal amplitude, namely one right-circular and one left-circular. The chiral molecules interact differently with the two circular components, giving them slightly different refractive indices and therefore slightly different propagation speeds inside the solution. After traversing the column, the two circular components are no longer at the same phase. When they recombine, the resultant is again plane-polarised, but the plane of polarisation has been rotated through an angle proportional to the path length and to the concentration of the chiral molecules.

(b) Specific rotation of a chiral solution is defined as the rotation per unit path length per unit concentration:

$$\alpha_0 = \frac{\alpha}{Lc} \text{ (from } \alpha = \alpha_0 Lc)$$

For convenience, L is conventionally expressed in decimetres ($1\text{dm} = 10\text{cm}$) and c in grams per cubic centimetre:

$$\alpha_0 = \frac{13°}{(1.0\text{dm})(0.20\text{g/cm}^3)} = 65°\,\text{dm}^{-1}(\text{g/cm}^3)^{-1}.$$

With these miscellaneous worked examples now complete, the threads of Physical Optics: wavefronts, coherence, interference, diffraction, and polarisation, should feel woven into a single picture. It is time to sharpen our thinking. Welcome to the Digging Deeper Exercise!

DIGGING DEEPER EXERCISE 8

EXERCISE 8A: BINDER QUESTIONS

Question 1

Two ordinary tungsten-filament lamps placed close together fail to produce a stable interference pattern on a nearby screen, no matter how carefully the geometry is set up. Account in physical terms for this observation.

Question 2

Account for the experimental observation that the spacing between consecutive dark fringes formed by an air wedge is the same whether the wedge is viewed close to its apex or far from it.

Question 3

A compact disc held under a fluorescent lamp shows iridescent colours which shift continuously as the disc is tilted relative to the lamp. Account in physical terms for the shifting of the colours.

Question 4

A student claims that, in a Young's double-slit experiment, the dark fringes are arbitrary regions of low intensity lying at no particular position relative to the bright fringes. Indicate, with reasons, whether the dark fringes occur at any specific positions relative to the bright fringes.

Question 5

Kipanga and Kipute argue about a Newton's-rings experiment. Kipanga claims that the centre of the ring pattern viewed in transmitted light is dark, just as it is in reflected light. Kipute claims that the centre viewed in transmitted light is bright. Resolve the disagreement.

Question 6

A student claims that diffraction and interference are two completely different physical phenomena, with diffraction occurring at slits and interference occurring between separate wave trains. Indicate, with reasons, whether the student's distinction is physically meaningful.

Question 7

A student watching a single-slit diffraction pattern claims that the central maximum is twice as wide as each of the secondary maxima, while a colleague claims it is the same width. Resolve the disagreement, giving the physical reason for the correct answer.

Question 8

Account in physical terms for the experimental fact that, when unpolarised light passes through a single Polaroid sheet, exactly half of the incident intensity is transmitted, regardless of the orientation of the Polaroid's transmission axis.

Question 9

Kipanga claims that Brewster's angle for a transparent medium is the same for every colour of light, since it depends only on the refractive index of the medium. Identify the error in this claim and account for the small dependence of Brewster's angle on wavelength.

Question 10

Show that the dark fringes in a Young's double-slit pattern occur exactly midway between consecutive bright fringes.

Question 11

A student claims that immersing a Young's double-slit apparatus in water should shift the entire interference pattern toward one side of the slits, because the wavelength has changed. Identify the error in this claim and account in physical terms for what actually happens to the pattern.

Question 12

Account in physical terms for the experimental observation that the radii of the dark rings in a Newton's-rings pattern increase as the square root of the ring number, rather than in direct proportion to the ring number.

Question 13

A thin soap film viewed in white light shows bright bands of colour, but a glass window-pane several millimetres thick viewed in the same white light shows no such colours at all. Account in physical terms for the difference.

EXERCISE 8B: REAL QUESTIONS

Question 14

In a Newton's-rings demonstration at a science fair in Morogoro, the plano-convex lens has been silvered on its lower surface, so that all the incident light is reflected at the lens itself and none reaches the flat glass plate below. Account in physical terms for the change, if any, that the silvering produces at the centre of the ring pattern observed in reflected light.

Question 15

A student wearing polarised sunglasses with a vertical transmission axis observes that the screen of her smartphone, viewed in landscape orientation under bright daylight, appears bright and clear when the phone is held one way, but turns almost completely black when the phone is rotated through ninety degrees about the viewing axis. Account in physical terms for this observation.

Question 16

A child blowing soap bubbles in a courtyard in Mtwara observes that each bubble shows a shifting pattern of rainbow colour bands that flow downward across the bubble's surface before the bubble bursts. Account in physical terms for the downward flow of the colour bands.

Question 17

A motorcyclist parking near a workshop in Mwanza notices, after a brief shower of rain, that a thin layer of petrol on a puddle shows brilliant colours that change as he tilts his head and changes his viewing angle. Account in physical terms for the change of colour with viewing angle.

Question 18

A fisherman on the northern shore of Lake Tanganyika wears polarised sunglasses with a vertical transmission axis to reduce the glare of midday sunlight reflected from the surface of the lake. Account in physical terms for the orientation of his line of sight relative to the lake surface at which the sunglasses are most effective.

Question 19

A driver travelling along the Morogoro–Iringa highway at midday observes a shimmering pool of "water" stretching across the tarmac far ahead of the vehicle, which appears to recede as he approaches. Account in physical terms for this observation.

Question 20

A spectacle-maker in Arusha coats a glass lens with a thin transparent layer of magnesium fluoride to reduce the intensity of light reflected at a wavelength of 550nm. Account in physical terms for the reduction in reflected intensity at that wavelength.

Question 21

A photographer in Iringa observes the clear midday sky through a polarising filter rotated in front of her camera lens. She notices that the sky directly overhead darkens visibly when the filter is rotated to one particular orientation, while the sky near the sun barely changes its brightness. Account in physical terms for the orientation-dependent darkening of the overhead sky.

Question 22

A holographic image embedded in the security label of an identification card has been recorded with red laser light. In a shop in Dar es Salaam, the image appears sharp and clearly three-dimensional under a red LED inspection lamp, but appears blurred and washed-out under the yellow sodium street-lamps outside. Account in physical terms for the difference.

Question 23

A student in Tanga places a clear plastic protractor between two crossed Polaroid sheets and observes a faint multicoloured pattern. When she squeezes the protractor gently between her fingers, the pattern brightens and becomes vivid; when she releases the squeeze, the pattern fades back to its faint state. Account in physical terms for the brightening of the pattern under squeezing.

Question 24

A trader at Kariakoo market in Dar es Salaam holds up a compact disc to a customer to demonstrate the spectrum of colours that appears on its surface in bright sunlight. The trader claims that the CD is *"splitting white light like a prism".* Account in physical terms for the distinction between the colour-spreading mechanism of the CD and that of a prism.

Question 25

A radiology technician at the Muhimbili National Hospital in Dar es Salaam is calibrating an ultrasound scanner. He notes that a 5MHz pulse maintains its frequency when crossing from a layer of subcutaneous fat into the underlying muscle, but that the wavelength inside the muscle is longer than inside the fat. Account in physical terms for this observation.

EXERCISE 8C: HOT QUESTIONS

Question 26

Light of wavelength 600nm in vacuum is incident at 50.0° on the surface of a glass slab of refractive index 1.50. Calculate the angle of refraction inside the glass, the wavelength of the light inside the glass, and the speed of the light inside the glass.

Question 27

A sodium lamp emits two close spectral lines whose mean wavelength is 589nm and whose frequency spread is 5.0×10^9Hz. Calculate the coherence time and the coherence length of the emitted light.

Question 28

In a Young's double-slit experiment, light of wavelength 550nm illuminates two slits separated by 0.25mm, and the fringes are observed on a screen 1.20m from the slits. Find the fringe width on the screen and the angular separation of consecutive bright fringes as seen from the slits.

Question 29

In a Young's double-slit experiment, light of wavelength 620nm produces fringes of spacing 1.55mm on a screen 1.50m from the slits. Determine the slit separation, and find the new fringe width if the slit separation is doubled while the wavelength and the screen distance are unchanged.

Question 30

In a Young's double-slit experiment with slit separation 0.45mm, screen distance 1.20m, and illumination of wavelength 589nm, a thin sheet of glass of refractive index 1.52 is placed in front of one of the slits and the central fringe shifts by 4.50mm from its original position. Calculate the thickness of the glass sheet.

Question 31

An air wedge is formed between two glass plates separated at one end by a strip of paper of thickness 0.040mm and length 8.0cm. Light of wavelength 589nm is incident normally on the wedge. Find the spacing between consecutive dark fringes, and the total number of dark fringes observed along the wedge.

Question 32

A thin film of refractive index 1.40 is illuminated normally by white light. Find the smallest film thickness that gives a bright reflected fringe at a wavelength of 480nm, and the next-smallest film thickness that gives a bright reflected fringe at the same wavelength.

Question 33

In a Newton's-rings experiment, light of wavelength 600nm is used and the plano-convex lens has a radius of curvature of 0.90m. Calculate the radius of the tenth dark ring measured from the centre, and find the radius of curvature of a lens that would give the same tenth-ring radius if the wavelength were changed to 500nm.

Question 34

In a Newton's-rings experiment, the diameter of the tenth dark ring is 5.00mm when there is air between the lens and the flat plate. When the gap is filled with a transparent liquid, the diameter of the same tenth ring shrinks to 4.30mm. Calculate the refractive index of the liquid.

Question 35

A monochromatic beam of wavelength 589nm passes through a single slit and forms a diffraction pattern on a screen 2.50m from the slit. The full width of the central maximum on the screen is 1.20cm. Find the width of the slit and the angular position of the third-order dark fringe.

Question 36

A compact disc has reflective tracks spaced 1.6 μm apart and is illuminated normally by white light. Find the diffraction angle of the first-order principal maximum for red light of wavelength 700nm and for blue light of wavelength 450nm.

SOUND, ELECTROMAGNETIC WAVES, AND THE DOPPLER EFFECT

INTRODUCTION

Jangwani Secondary School in Dar es Salaam, mid-morning on a Wednesday. The Inter-School Physics Club is in session. Mr. Akilikubwa has brought Kipute and Kipanga from Miono Secondary School on a three-day study tour, and the Jangwani host teacher has loaned them a Form Five classroom that faces United Nations Road, the busy corridor between Kariakoo and Muhimbili National Hospital. The windows are open because the fan is broken and the morning is already 32 degrees. Outside, traffic crawls. Inside, the room is quiet.

Monalisa, a Form Five girl from Jangwani who lives in Upanga and has heard ten thousand sirens in her life, is at the blackboard finishing a diagram of a stationary wave. **Khadija**, a Form Six girl from Jangwani who reads physics textbooks for entertainment, is correcting a sign in one of her equations with the calm precision of a someone who has waited four years for this exact moment. **Kipanga** is staring at the diagram and pretending to understand. **Kipute** is checking Khadija's correction and quietly noting that she is, on this occasion, correct. **Mr. Akilikubwa** is at the back of the room, watching the four of them with the expression of a fisherman who has just seen his float twitch.

From the road outside, a sound begins.

NEEEEEEEEEEEE—

An ambulance, racing from Kariakoo north to Muhimbili. It is coming fast. The class falls completely silent. Khadija's chalk freezes against the board. Kipanga stops pretending to understand. The siren reaches the school gate at maximum loudness, the most arresting noise a human throat or an electronic horn can make.

NEEEEEEEEEEEEE—

And then it passes, accelerating north towards the hospital, and the sound changes.

nooooooo—

Same siren. Same machine. Same driver. Same morning. A completely different note.

Kipanga, who has been silent since arriving at Jangwani because the city classroom intimidates him slightly, finally finds a subject he has authority on, namely sirens.

Kipanga: *"Sir, this one is easy. The driver pressed the siren button hard while coming, and soft while going. At Miono we have the same thing with the thief-alarm whistles at the school shop. A hard press is louder and a soft press is quieter."*

Monalisa turns from the blackboard. Her eyebrow rises with the slow patience of someone who has explained Dar to visitors many times before.

Monalisa: *"Kipanga, that is an electronic siren on a Land Cruiser ambulance. It does not have a press-button. It has one switch, on or off, and the people inside are not playing it like a piano. They are too busy keeping the patient alive. Also, in the four years I have lived next to this road, I have not once heard an ambulance pass slowly. They come fast, they leave fast, and the note still changes."*

Kipanga is not yet ready to surrender.

Kipanga: *"Then perhaps the driver opened the window while coming and closed it while going?"*

Khadija puts down her chalk and joins in, with the composed confidence of a Form Six student who has read about this before and now intends to demonstrate it.

Khadija: *"Kipanga, it is the air itself. The sound waves bunch together in front of the ambulance, the way daladalas pile up at the Mwenge roundabout in the evening rush, and they stretch apart behind it, the way the same daladalas look on the Bagamoyo Road at midnight. Bunched-up waves are short. Short waves are high in pitch. Stretched-out waves are long. Long waves are low in pitch. There. The driver and the windows have nothing to do with it."*

Kipute, who has been resting her chin in her hand while thinking, now drops her hand and taps Khadija on the elbow.

Kipute: *"Khadija, your daladalas-at-Mwenge picture is mostly right. But why does the ambulance bunch the waves up? The siren is producing the same number of beeps per second whether the ambulance is parked, or moving slowly, or racing to Muhimbili. The frequency at the source is fixed. So what changes?"*

Khadija opens his mouth, closes it, opens it again.

Khadija: *"Because the source itself moves between one beep and the next?"*

Kipute: *"Yes. But say more."*

Khadija: "I... will think about it."

Mr. Akilikubwa, who has been watching this with the expression of a man whose float has now disappeared under the surface, walks slowly to the blackboard. He erases Monalisa's stationary wave with one sweep of the duster. Monalisa winces. He writes, in chalk, a single equation.

$$f' = f \times \frac{v \pm u_o}{v \mp u_s}$$

He turns to face the class.

Mr. Akilikubwa: *"This equation is the entire ambulance. Everything Kipanga said is wrong, and we love him for trying. Everything Khadija said is half right, which is the most dangerous kind of right because the other half is missing. Everything Kipute is about to say is correct, and that is why we sent her on this study tour. The siren did nothing. The driver did nothing. The universe did one very specific thing. Motion bends perception before it bends sound. This chapter is the unpacking of that sentence."*

Monalisa, recovered from the loss of her stationary wave.

Monalisa: *"Sir, but why is there a plus on top and a minus on the bottom? And why are the signs paired like this instead of being the other way round?"*

Mr. Akilikubwa: *"That, is what the next seventy pages are for. By the time we finish this chapter, you will know not only why the formula looks the way it does, but also why the same kind of formula explains the colour of a distant galaxy, the speed of a car caught by a speed camera on the Morogoro Road, the way a bat finds a moth in the dark inside Kondoa cave, and why your TBC radio sometimes picks up broadcasts from Mozambican radio stations at night but never at noon. All of it, every one of those phenomena, follows from one ambulance, one road, and one equation."*

Kipanga, in a whisper to Kipute.

Kipanga: *"Kipute, I think I should pay much more attention in Form Five than I did in Form Four."*

Kipute: *"Kipanga, you are in Form Five."*

Welcome to Chapter 9. The previous three chapters built the toolkit that this chapter spends.

This chapter takes those three toolkits and applies them to two specific media that students hear and see every day, namely sound waves in air and electromagnetic waves in vacuum. It then introduces one phenomenon, the Doppler effect, which acts on both media in essentially the same way and which turns out to be the single most useful idea in classical physics for measuring motion at a distance.

Two debts from Chapter 7 are paid in this chapter, and you should know which ones before we begin.

The first debt is the speed of a wave in a gas. Chapter 7 derived the wave speed for a stretched string, a solid rod, and a liquid, all of them from the same square-root rule that the speed equals the root of an elastic property over an inertial property of the medium. Gases were postponed at the time, because the elastic property of a gas depends on whether the compression at each cycle of the wave is isothermal (slow enough for heat to flow) or adiabatic (too fast for heat to flow), and Chapter 7 had not yet earned the right to bring in the adiabatic process from Chapter 5. We have done the homework since. The section that follows on the speed of sound applies the square-root rule to a gas, identifies the adiabatic bulk modulus as the elastic property that fits, and pays the debt in full.

The second debt is the harmonic spectrum of an air column. Chapter 7 derived the harmonic spectrum of a stretched string fixed at both ends, and then promised that the same boundary-condition reasoning would deliver the harmonic spectrum of an air column in a pipe. The pipe argument was postponed on the grounds that pipes are really about sound, and that sound deserved its own chapter. The section on resonance in pipes that comes immediately after, honours that promise.

Let us begin.

THE NATURE OF SOUND

The introduction has done the dramatic work for us. We have an ambulance, a road, a classroom at Jangwani with the windows open because the fan is broken, and Mr. Akilikubwa's equation sitting on the blackboard like a man waiting for a daladala. Before we can ask why the siren rose in pitch coming towards the school and dropped in pitch going away, we must first agree, calmly and without sirens, on what a sound actually is. This first section is the calm one. By its end you will know three things, and you will know them firmly enough that the rest of the chapter can rest its weight on them without trembling. The first is what kind of wave a sound is. The second is what the ear hears when a sound arrives, and what each piece of that hearing corresponds to in the wave itself. The third is why a diver paddling five metres below the surface at Kunduchi Beach cannot hear a friend in the boat above him shouting his name. No equation in this section will frighten you. The serious calculations begin in the next section, when the speed of sound enters the chapter and refuses to leave.

What kind of wave is a sound, really?

The biggest fact about sound is the one we already know from Chapter 7, but we will state it once more because the rest of the chapter swings off it. A sound is a **longitudinal mechanical wave. Longitudinal** *because the particles of the medium oscillate in the same direction in which the wave is travelling, not perpendicular to it.* **Mechanical** *because it needs a medium: there is no sound in a vacuum,* a point that Mr. Akilikubwa will return to several times before the chapter is finished, often at exam time.

What Chapter 9 adds to Chapter 7 is the medium-specific detail. The medium that carries sound in everyday life is **a fluid**, almost always **a gas** (the air around us), sometimes **a liquid** (the seawater the diver is swimming through), and on **rare** interesting occasions **a solid** (the steel rail along the TAZARA line that the trackside worker presses his ear to). In every one of those media the disturbance is the same: a small local fluctuation in pressure and density. A vibrating source pushes the air immediately around it slightly inwards; the air recoils outwards when the push is released; the disturbance propagates outwards through the medium as alternating regions of slightly increased pressure (called **compressions**) and slightly decreased pressure (called **rarefactions**); each particle of the medium, on average, ends up exactly where it started. Energy travels outwards. Air does not. The diagram below shows exactly this picture.

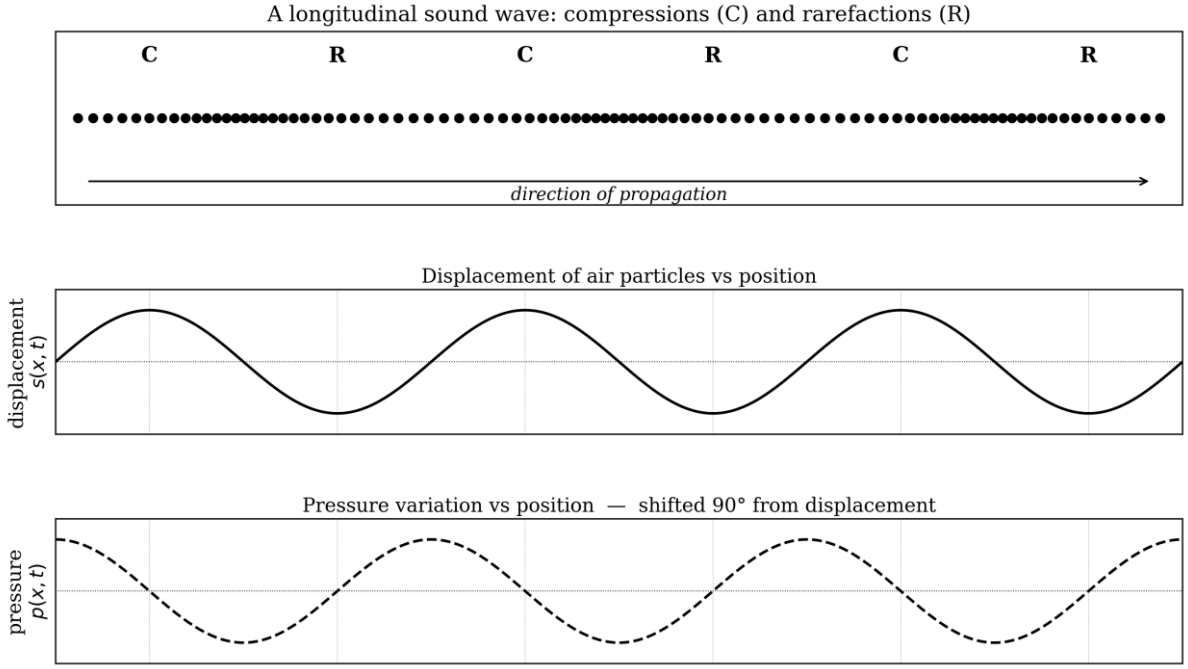


Figure: ***A longitudinal sound wave.*** *Top: the air molecules cluster (compressions, marked C) and spread out (rarefactions, marked R) as the wave moves along. Middle: displacement of each particle from its rest position. Bottom: pressure variation. The two curves are sinusoids of the same wavelength, but pressure peaks fall exactly where displacement is zero.*

Two graphs sit next to every textbook picture of a sound wave, and the relationship between them is the source of more confusion in this chapter than any other piece of geometry. The first graph shows the **displacement** of each particle from its rest position. The second graph shows the **pressure variation** (the small departure of the pressure from atmospheric) at the same instant, as a function of the same position. Both graphs are sinusoidal. Both have the same wavelength. Both travel at the same speed. But they are **ninety degrees out of phase with each other,** and they are out of phase for a reason that becomes obvious the moment you stop to think about it.

Consider the place along the wave where two groups of particles are converging from opposite sides. The particles on the left are moving rightwards; the particles on the right are moving leftwards. Right at the spot where the two groups meet, the air is being squeezed from both sides at once. The pressure there is at its **maximum**. But the particles at that exact spot have zero net displacement: equal amounts of particle have come in from each side, the local mean position has not moved. So at the point of maximum pressure, the displacement is zero (a **node** of displacement). Where the air has bulged most strongly to one side, the pressure is back to atmospheric (a node of pressure), but the displacement is at its maximum there (an **antinode** of displacement). The two patterns interlock like the teeth of a zipper. Pressure peaks fall exactly where displacement is zero, and the other way around. This is what "ninety degrees out of phase" means in pictures, and it is the geometry that the resonance-tube experiment in the next-but-one section lives on. If you can hold this picture clearly in your head, you will save yourself a lot of suffering when we come to the closed-pipe boundary condition.

Mathematically, the displacement of a particle at position x at time t can be written, exactly as in Chapter 7:

$$s(x,t) = s_{max}\sin(\omega t - kx)$$

The pressure variation at the same point at the same time is then:

$$p(x,t) = p_{max}\cos(\omega t - kx)$$

A cosine is a sine shifted by a quarter of a cycle, and that quarter-cycle shift is exactly the ninety-degree offset we have just argued for. The two amplitudes are not independent. They are linked by the elastic property of the medium, and the link is worth deriving once because it shows the wave-speed formula from Chapter 7 in action.

Consider a slab of air whose rest position stretches from x to x + Δx, with cross-sectional area A. When the wave passes through, the left face moves to x + s(x, t) and the right face moves to x + Δx + s(x + Δx, t). The slab's new length is Δx + s(x + Δx, t) − s(x, t). The fractional change in length, which (because the slab is thin) equals the fractional change in volume, is obtained by dividing by Δx and letting Δx → 0:

$$\frac{\Delta V}{V} = \frac{\partial s}{\partial x}$$

Differentiate the wave. From the wave expression for s(x, t) above:

$$\frac{\partial s}{\partial x} = -k\,s_{max}\cos(\omega t - kx)$$

Apply the bulk modulus:
The bulk modulus K of the medium, defined in Chapter 5, says that a fractional volume change produces a pressure change in the opposite sense (a compression raises pressure, a rarefaction lowers it):

$$p = -K\frac{\Delta V}{V} = -K\frac{\partial s}{\partial x} = K\,k\,s_{max}\cos(\omega t - kx)$$

Comparing with the form $p = p_{max}\cos(\omega t - kx)$ we wrote earlier, the pressure amplitude is therefore:

$$p_{max} = K\,k\,s_{max}$$

Use the wave-speed formula:
The unified formula from Chapter 7 gives the wave speed in a fluid as $v = \sqrt{(K/\rho)}$, from which $K = \rho v^2$. The wave number is $k = \omega/v$. Substituting both into the line above:

$$p_{max} = (\rho v^2)\left(\frac{\omega}{v}\right)s_{max}$$

The factor of v in the numerator and denominator partially cancels; the surviving v carries the whole imprint of the elastic property of the medium. The result is:

$$\mathbf{p_{max} = \rho v \omega s_{max}}$$

where ρ is the density of the medium, v is the speed of sound in the medium, and $\omega = 2\pi f$ is the angular frequency.

Pitch, loudness, and quality: the three things the ear hears

A musician sitting at the front of a hall on Sokoine Avenue can listen to a single note from a single instrument and describe three completely different things about it.

The first *is how high or low the note is.* A piccolo is high; a tuba is low. This perceptual property is called the **pitch** of the note.

The second *is how loud the note is.* A whisper is quiet; a market trader at Kariakoo announcing the price of mangoes is loud. This perceptual property is called the **loudness**.

The third *is what kind of instrument is making the note,* even when the pitch and loudness of two notes are matched. A clarinet and a guitar playing the same note at the same loudness do not sound the same; the listener can tell them apart instantly. This perceptual property is called the **quality** of the note, also called the ***timbre***.

Each of these three perceptual properties corresponds to one specific feature of the wave that the ear receives. The correspondence is not vague; it is one-to-one, and it is worth memorising once and then using freely for the rest of the chapter.

Pitch corresponds to **frequency**. A high-pitched note is a wave of high frequency; a low-pitched note is a wave of low frequency. Doubling the frequency raises the pitch by exactly one **octave**, which is the musical interval between middle C and the next C above it. The piano A above middle C has a frequency of 440Hz; A two octaves above that is 1760Hz; A two octaves below is 110Hz. Two notes sound musically equivalent (same letter name, just at different heights) when their frequencies are in a simple ratio.

Loudness corresponds to **intensity**. The intensity of a wave was defined in Chapter 7 as the power carried across a unit area per second, and intensity scales as the square of the amplitude. A louder note is a wave of larger pressure amplitude (or equivalently larger displacement amplitude). The detailed mathematics of the loudness scale, including the logarithmic decibel scale that the rest of the world uses to write about loudness, lives a few sections from now in the section on the decibel.

Quality corresponds to the **harmonic spectrum**. A real musical note from a real instrument is almost never a single sinusoid at a single frequency. It is a sum of many sinusoids: a strongest one at the fundamental frequency f_1, and weaker ones at integer multiples of f_1, called the second harmonic at $2f_1$, the third harmonic at $3f_1$, and so on. Chapter 7 met this idea when it discussed the vibrating string. The exact relative amplitudes of the higher harmonics depend on how the source is excited: a

string plucked at its middle gives a different harmonic mix from the same string plucked near the bridge, a clarinet's column of air gives a different mix from a guitar's string, and a human voice singing the vowel AAH gives a different mix from the same voice singing EEH. The ear takes the full mix in and presents it to the brain as a single percept called timbre. Same fundamental, different harmonic recipe, different instrument: that is the whole of quality.

A small but important consequence of this. When you hear your own voice from a recording for the first time, you flinch, because the recording does not sound like the voice in your head. The voice in your head is the part your skull bones conduct directly to your inner ear, which favours the lower harmonics. The voice in the recording is what a microphone in front of your mouth heard, which favours the higher harmonics carried through the air. Same fundamental, different harmonic mix, different timbre, and so a different perceived voice. This is not your phone playing tricks on you. It is the harmonic spectrum doing its accounting.

The audible range, and the worlds on either side of it

The human ear is sensitive to wave frequencies between roughly **20Hz at the low end and 20000Hz (20kHz) at the high end**, in a healthy young person. The exact bounds shift with age: the upper bound retreats first, so that an adult of fifty rarely hears above 12kHz, a finding that hi-fi shop assistants in Kariakoo have learned to use to their advantage. The full span of two-and-a-half orders of magnitude is called the **audible range**. Below the lower bound and above the upper bound, the disturbance is still a perfectly good sound wave; the human ear simply does not register it.

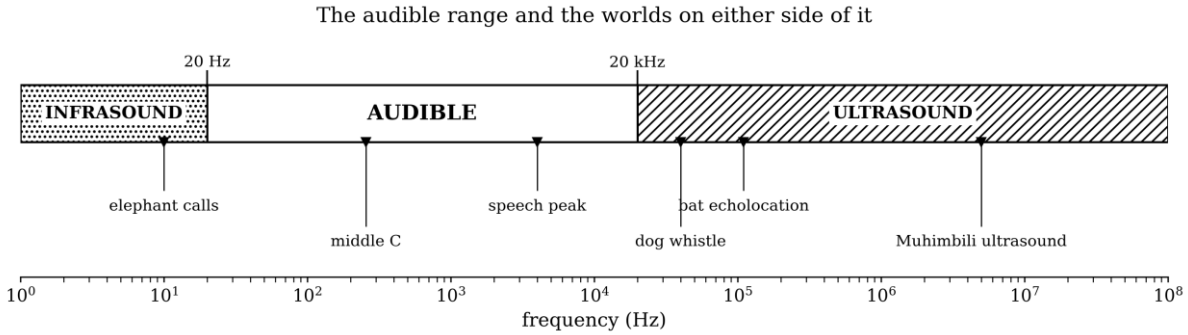


Figure: *The audible range on a logarithmic frequency axis, with neighbouring infrasound and ultrasound. The 50Hz tick is the hum a Tanzanian mains transformer makes when the wiring is sloppy.*

Below 20Hz the vibrations are called **infrasound**. Earthquakes produce them; large explosions produce them; elephants use them to communicate over kilometres of savannah, which is one of the more surprising practical facts of biology. The vibrations are inaudible but felt: a powerful infrasonic pulse can make a window rattle without making any audible noise.

Above 20kHz the vibrations are called **ultrasound**. Dogs can hear up to about 45kHz, which is why a dog whistle blown at 35kHz drives a dog mad and a human does not even know the whistle has been blown. Bats hunt at frequencies up to 110kHz and read the echoes of their own ultrasonic pulses to navigate; this is the topic of a Doppler example a few sections from now. Medical ultrasound scanners at Muhimbili operate between 2MHz and 15MHz, three orders of magnitude above audibility, because the wavelength at those frequencies is short enough to resolve organs of a few millimetres. None of these waves are physically different from the sound of a school bell. They are simply at frequencies the human ear cannot follow.

Why the diver cannot hear the friend in the boat

A swimmer at Kunduchi Beach, paddling at the surface, hears the boat engine perfectly. A diver swimming five metres below the surface hears nothing, even when a friend in the boat is shouting at the top of his lungs. The diver's ear works perfectly. The friend's lungs are at full pressure. *So why does no sound reach the diver?*

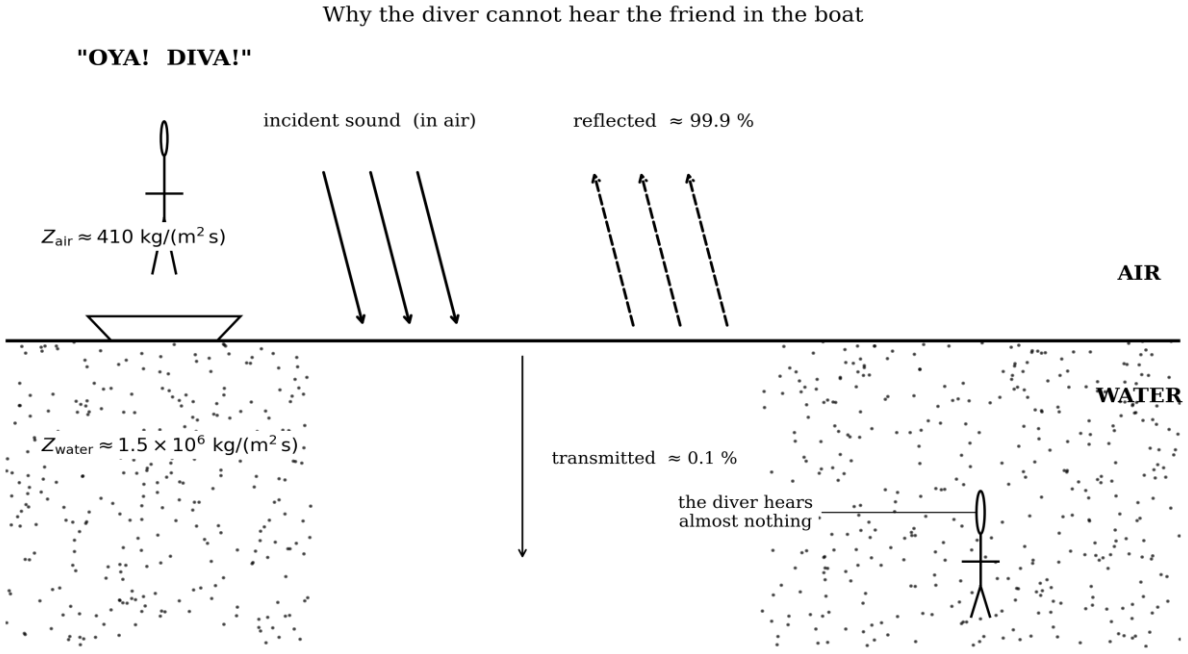


Figure: *A loud shout in air carries less than 0.1% of its power into water. The other 99.9% is reflected back at the surface. The diver swims, in practice, in silence.*

The answer is one of the most useful pieces of physics in the chapter, and it is the first place where Chapter 9 pays back its debt to Chapter 7's discussion of the wave-speed formula. The key idea is **acoustic impedance**. When a sound wave travelling in one medium meets the boundary of a second medium, two things can happen. Part of the wave *passes into the second medium* (**transmission**); part of the wave *bounces back into the first medium* (**reflection**). The fraction that gets through depends on how different the two media are, in a very precise sense. Each medium has an **acoustic impedance Z**, defined by:

$$\mathbf{Z = \rho v}$$

the product of the medium's density and its speed of sound. For air at room temperature, $\rho \approx 1.2$kg/m³ and $v \approx 340$m/s, giving $Z_{air} \approx 410$kg/(m²s). For seawater at twenty degrees Celsius, $\rho \approx 1030$kg/m³ and $v \approx 1500$m/s, giving $Z_{water} \approx 1.5\times10^6$ kg/(m²s). The two impedances differ by a factor of about 3700, which is the kind of ratio that sets off alarm bells in every wave physicist who has ever stared at a boundary problem.

The fraction of sound power transmitted across the boundary at normal incidence is given by:

$$T = \frac{4\,Z_{air}\,Z_{water}}{(Z_{air} + Z_{water})^2}$$

Substituting our two numbers gives $T \approx (4 \times 410 \times 1.5\times10^6) / (1.5\times10^6)^2 \approx 1.1\times10^{-3}$. About 0.1% of the power in the shout makes it through the water surface; the other 99.9% is reflected back into the air. By the time the shout has spread spherically through five metres of water as well, the intensity at the diver's ear is well below the threshold of hearing. The shout is, for the diver, silent.

This is called the **impedance-mismatch argument**, and it is everywhere in physics whenever a wave crosses a boundary between two media of very different stiffness or density. We will meet it again in the applications section when we discuss why an ultrasound scanner needs a gel layer between the transducer and the patient's skin: the gel matches the impedance between the transducer and the body well enough that almost all the ultrasound energy enters the body, instead of bouncing back off the skin like an air-shout off the surface of the sea.

With the picture now firmly in place, the worked examples can begin.

BINDER Example 1

A loudspeaker in the studio of Clouds FM in Kinondoni emits a steady note at 256Hz. The pressure amplitude of the wave one metre in front of the speaker is 0.20Pa. Take the speed of sound in air as 343m/s and the density of air as 1.2kg/m³.

(a) Find the wavelength of the wave.

(b) Find the displacement amplitude of an air particle one metre in front of the speaker.

Solution

(a) The wavelength is the speed of sound divided by the frequency:

$$\lambda = \frac{v}{f} = \frac{343\text{m/s}}{256\text{Hz}} = 1.34\text{m}$$

(b) The displacement amplitude is connected to the pressure amplitude by:

$$s_{max} = \frac{p_{max}}{\rho v \times 2\pi f}$$

Substituting the numerical values:

$$s_{max} = \frac{0.20\text{Pa}}{(1.2\ \text{kg/m}^3)(343\text{m/s})(2\pi)(256\ \text{Hz})} \approx 3.0 \times 10^{-7}\text{m} \approx 0.30\mu\text{m}$$

Making sense of the answer: *The wavelength of 1.34m is comparable to a child's height, which is the reason a 256Hz note bends easily round door frames and corners of buildings: a wavelength of about a metre diffracts strongly around obstacles a metre or two in size. The displacement amplitude of 0.30 micrometres is far smaller than the diameter of a human hair (which is about 80 micrometres), and it is one of the central facts about hearing: the ear can detect particle displacements that are smaller than a wavelength of visible light.*

Think like a physicist: *Two amplitudes describe the same wave, and they are connected by a single equation. The pressure amplitude p_{max} is what an engineer designing a loudspeaker measures with a microphone. The displacement amplitude s_{max} is what a physicist analysing the air molecules computes from the equation of motion. They are the same physical wave seen through two different mathematical windows. Whenever a problem gives you one and demands the other, the bridge between them is always $p_{max} = \rho v \omega s_{max}$.*

REAL Example 2

A dog whistle in the hand of a frustrated dog-owner in Wete emits a steady note at 40kHz. A human a few metres away hears absolutely nothing; the family dog, sitting equally near, looks up sharply and starts whining about an injustice he cannot quite name.

(a) Explain why the dog hears the whistle and the human does not.
(b) Find the wavelength of the whistle's note in air at 343m/s.
(c) A fruit bat flying near the house uses echolocation at frequencies up to 110kHz. Would the bat hear the same whistle as a low note or a high note compared with its own calls?
(d) Why would the same whistle, applied a few centimetres from the dog's ear, hurt the dog rather than charm him?

Solution

(a) The human auditory range tops out at about 20kHz in a young adult and lower in an older one. 40kHz lies well beyond the human upper cut-off; the human eardrum cannot respond fast enough, so the wave passes through the room undetected by the human ear. The dog auditory range extends up to about 45kHz, so 40kHz lies inside what the dog can hear. The whistle is loud, but the loudness is delivered in a frequency band that only one of the two species has equipment for.

(b) The wavelength in air:

$$\lambda = \frac{v}{f} = \frac{343\text{m/s}}{40000\text{Hz}} = 8.6 \times 10^{-3}\text{m} = 8.6\text{mm}$$

(c) The bat's echolocation calls reach up to 110kHz. The whistle at 40kHz lies near the bottom of the bat's hearing range. For the bat, the whistle would sound like a low note, just as a 200Hz human voice sounds low compared with a 1000Hz whistle to a human listener. Frequencies are perceived as high or low relative to the upper end of one's own audible range, and the bat's upper end is much higher than the human's.

(d) Loudness has nothing directly to do with frequency. A 40kHz ultrasonic note can carry enormous pressure amplitude and intensity. The dog's eardrum, although it can follow the frequency, cannot withstand arbitrary intensity. A close-range blast delivers a large pressure amplitude through a wave whose frequency happens to be inside the dog's hearing band, and the eardrum responds painfully to the intensity. The same intensity at 1000Hz, well inside both species' hearing, would hurt a human ear just as much. Frequency makes a sound audible to a given species; intensity makes it painful.

Making sense of the answer: *An 8.6mm wavelength is about the width of a small fingernail. Such short wavelengths reflect efficiently off small surfaces, which is exactly why bats can resolve insects a few millimetres in size by echolocation; the same short wavelengths diffract poorly round large obstacles, which is why the ultrasonic part of the whistle does not travel well behind furniture or round corners, while the lower-frequency hisses produced as side-effects of the blowing do. This same wavelength argument is why an ultrasound transducer at Muhimbili operating at 5MHz, with a wavelength in tissue of about 0.3mm, can resolve features one or two millimetres across; an audible-frequency probe with a wavelength of one full metre would resolve nothing finer than the patient's leg.*

Think like a physicist: *Audibility is a species-specific window onto a universal physics. There is no special category of "sound that humans cannot hear", and no special category of "sound that dogs can hear"; there is only the physical wave, with its frequency, and there are different ears with different upper cut-offs. The same physical wave that is silent to a human and audible to a dog is audible-but-low to a bat and audible-but-high to a baby (whose audible range typically extends slightly above an adult's). When a question talks about hearing, always separate the physics of the wave from the biology of the listener; they live in different sentences.*

With this section behind us, three perceptual properties (pitch, loudness, quality) have been mapped cleanly onto three wave properties (frequency, intensity, harmonic spectrum). One useful boundary trick (acoustic impedance mismatch) has explained why most sound waves bounce back at most boundaries; we will use the same trick again at the resonance tube and again at the medical-ultrasound transducer. We have not yet, however, addressed the single most-asked question of any chapter on sound: *how fast does it travel?* Chapter 7 promised the answer would arrive in this chapter, and the moment has come to deliver. The next section unwraps the speed of sound: what Newton calculated, why his answer was wrong by sixteen per cent, and what Laplace did about it.

SPEED OF SOUND: NEWTON, LAPLACE, AND TEMPERATURE

The previous section settled the perceptual side of sound. We know what a sound is, what the ear does with it, and why one species' silence is another species' alarm. But every one of those statements quietly assumed something we have not yet computed: the speed at which a sound wave actually travels through the medium. Chapter 7 promised us this speed and then walked away whistling, on the grounds that gases required an adiabatic argument that Chapter 7 had not yet earned the right to make. The argument has since been earned in Chapter 5, where the adiabatic process was derived in full. The promise comes due here.

The story has three protagonists. Sir Isaac Newton, who in 1687 wrote down a formula for the speed of sound in a gas and was wrong by sixteen per cent. Pierre-Simon Laplace, who in 1816 looked at the same formula, identified one missing piece, slotted it in, and made the formula right to the third significant figure. And the speed of sound itself, which has the inconvenient habit of caring strongly about temperature but not at all about pressure, an asymmetry that catches the unwary on examination day.

The unified wave-speed formula meets a gas

From Chapter 7 we already know that, in every elastic medium, the speed of a mechanical wave is given by the square root of an elastic property of the medium divided by an inertial property of the medium:

$$v = \sqrt{\frac{\text{elastic property}}{\text{inertial property}}}$$

For a gas the inertial property is straightforward: it is the volume mass density ρ. That part needs no debate. The elastic property is the part where Newton and Laplace will disagree, and it is the **bulk modulus** of the gas. Recall from Chapter 5 that the bulk modulus is the pressure required to produce a unit fractional decrease in volume, $K = -V(dP/dV)$. The trouble is that a gas can have two different bulk moduli depending on whether the compression that defines dV is carried out slowly (with enough time for heat to flow in or out and keep the temperature equal to the surroundings) or quickly (so fast that no appreciable heat flows during the compression). One number does not fit both situations, and the choice between them is our next task.

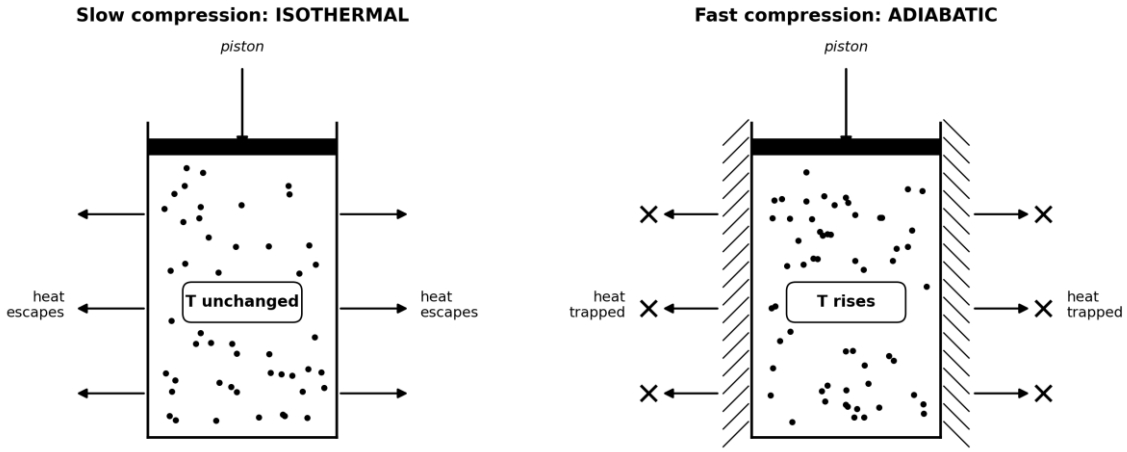


Figure: ***A gas in a cylinder being pushed by a piston in two limits.*** *On the left, the piston moves slowly enough that heat flows freely through the walls; the temperature of the gas stays equal to the surroundings; the compression is isothermal. On the right, the piston moves fast enough that no heat escapes during the compression (the arrows with crosses indicate that heat transfer is blocked or absent); the temperature of the gas rises; the compression is adiabatic.*

Newton's first attempt: isothermal compression

Newton's instinct in 1687 was simple. As the wave passes through the air, the air at each location alternately compresses and rarefies. Newton assumed the compression was slow enough that the air at each location stayed at the same temperature as the air around it: an **isothermal** compression. For an ideal gas held at constant temperature, the equation of state $PV = nRT$ reduces to $PV = \text{constant}$, from which a short calculation gives the isothermal bulk modulus:

$$K_T = -V\frac{dP}{dV} = P$$

So Newton substituted P for the elastic property:

$$v_{\text{Newton}} = \sqrt{\frac{P}{\rho}}$$

Plug in the values for air at standard conditions: pressure $P = 1.013 \times 10^5 \text{Pa}$, density $\rho = 1.29 \text{kg/m}^3$. The arithmetic gives:

$$v_{\text{Newton}} = \sqrt{\frac{1.013 \times 10^5 \text{Pa}}{1.29 \text{kg/m}^3}} \approx 280 \text{m/s}$$

Newton then went to find experimentalists to confirm the figure.

The sixteen-per-cent embarrassment

The experimentalists came back with a different number. Careful measurements of the time taken for a gun's report to reach an observer at a known distance gave the speed of sound in air at $0°\text{C}$ as roughly 332m/s, not 280m/s. The discrepancy is too large to be experimental error. Newton's calculation is off by:

$$\frac{332 - 280}{332} \times 100\% \approx 16\%$$

Sixteen per cent is the gap between physics and almost-physics. Something in the model is physically wrong, and the model has to be repaired. For more than a hundred years no one quite knew what to repair. The repair, when it arrived, did not change a single number in the formula; it changed the interpretation of one of the symbols. That is the kind of fix every physicist envies.

Laplace's correction: adiabatic compression

Laplace's 1816 insight was this. Sound is not a slow phenomenon. The audible frequencies start at 20Hz, which means each compression-and-rarefaction cycle takes at most 50ms; at 1kHz each cycle takes one millisecond; at 20kHz each cycle takes fifty microseconds. Heat conduction through air at this speed is far too slow to do anything useful. During a single cycle of the wave, the gas at any given location is compressed and then rarefied without any appreciable heat being able to leave or enter. The compression is therefore **adiabatic**, not isothermal. The right elastic property is not the isothermal bulk modulus but the **adiabatic bulk modulus**, derived in Chapter 5:

$$K_S = \gamma P$$

where $\gamma = C_p/C_v$ is the ratio of the specific heat capacities of the gas at constant pressure and constant volume. For diatomic gases such as air (mostly nitrogen and oxygen, both diatomic), $\gamma = 1.40$:

$$v = \sqrt{\frac{\gamma P}{\rho}}$$

This is the **Newton–Laplace formula**. Substituting $\gamma = 1.40$, $P = 1.013 \times 10^5 \text{Pa}$, $\rho = 1.29 \text{kg/m}^3$:

$$v = \sqrt{\frac{(1.40)(1.013 \times 10^5 \text{Pa})}{1.29 \text{kg/m}^3}} \approx 331 \text{m/s}$$

This agrees with the measured value to better than half a per cent. Laplace's correction is the difference between physics and almost-physics, and the whole of it lives in one symbol: the factor of γ inside the square root.

Temperature: the dependence the formula does have

Examiners love a question that asks for the speed of sound at a temperature that is not standard. The formula above is in terms of pressure and density, and both of those change with temperature. There is a cleaner form in which only the temperature appears. Substitute the ideal-gas law in the convenient form $P = \rho RT/M$ (where $R = 8.314 \text{J/(molK)}$ is the universal gas constant and M is the molar mass of the gas) into the Newton–Laplace formula:

$$\mathbf{v} = \sqrt{\frac{\gamma P}{\rho}} = \sqrt{\frac{\gamma \rho RT/M}{\rho}} = \sqrt{\boldsymbol{\frac{\gamma RT}{M}}}$$

The density cancels exactly. The speed of sound in an ideal gas depends only on the absolute temperature and on the molar mass of the gas. In particular, for any fixed gas (so γ and M are fixed):

$$\mathbf{v \propto \sqrt{T}}$$

This is the most important single line in this section. Two important consequences:

1) Air at 0°C (273K) gives $v = 331 m/s$; air at 25°C (298K) gives $v = 331 \times \sqrt{(298/273)} \approx 346 m/s$. An engineering rule of thumb worth memorising is the linear approximation $v \approx 331 + 0.6T_C$, with T_C in Celsius; the rule is good to better than one per cent over the everyday range $\mathbf{0–40°C}$.
2) The temperature at which v doubles compared with its value at 0°C satisfies $\sqrt{(T/273)} = 2$, giving $T = 4 \times 273K = 1092K = 819°C$. You will not encounter that temperature when whistling on the school field, but a furnace at a cement plant at Tanga reaches it without difficulty.

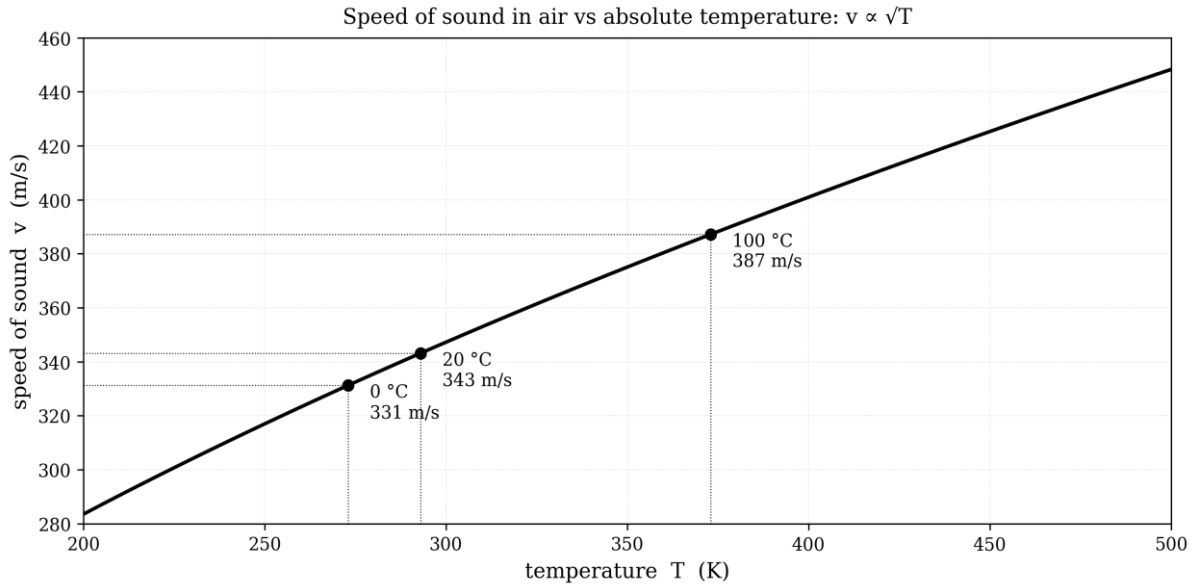


Figure: ***Speed of sound in air vs absolute temperature.*** *The curve is a square root of T, anchored at 273K → 331m/s. Doubling v means quadrupling T.*

Pressure: the dependence the formula does NOT have

From $v = \sqrt{(\gamma P/\rho)}$ it certainly looks as though v should depend on P. At first glance, halving the pressure should multiply the speed by $1/\sqrt{2}$, and doubling the pressure should multiply it by $\sqrt{2}$. But this is a trick the formula plays on the careless reader. At a **fixed temperature**, the ideal-gas law gives $\rho \propto P$: *doubling the pressure* of a fixed mass of gas at fixed temperature also *doubles its density*. The ratio P/ρ therefore stays exactly constant. The speed does not change.

This is why the speed of sound at the summit of Mount Kilimanjaro, where the air pressure is roughly half its sea-level value, is the same as the speed of sound on the beach at Dar es Salaam at the same temperature. The **lower** pressure on Kilimanjaro is exactly compensated by the **lower** air density there. Only the temperature, which is about 30°C lower on Kilimanjaro, gives the climber a slower sound; and that slowing is about 5%, much less than a naive use of the pressure ratio would predict. The asymmetry, in which temperature matters while pressure does not, is one of those facts the student must internalise and never write down in reverse on an examination paper.

Humidity: a small but real correction

Humid air is slightly less dense than dry air at the same temperature and pressure, because water vapour ($M = 18 g/mol$) replaces nitrogen ($M = 28 g/mol$) and oxygen ($M = 32 g/mol$), both of which are heavier. From $v = \sqrt{(\gamma RT/M)}$, smaller M means larger v. In Dar es Salaam at 30°C and 80% relative humidity, the speed of sound is about 0.6% higher than in completely dry air at the same temperature. This is the reason a singer's voice carries a fraction further on a humid coastal morning than in dry Dodoma highland air at the same thermometer reading. The correction is usually too small to matter in calculations, but it deserves to appear at least once in your notes so that it never catches you by surprise later.

Before the ideas have a chance to drift away, let us anchor them with three worked examples.

BINDER Example 3

Calculate the speed of sound in dry air at 20°C using the Newton–Laplace formula, taking $\gamma = 1.40$, molar mass of dry air $28.96 g/mol$ and $R = 8.314 J/(mol \cdot K)$. Cross-check the answer using the engineering rule $v \approx 331 + 0.6T_C$.

Solution

Convert the temperature: $T = 20 + 273 = 293K$.
Substitute into the formula:

$$v = \sqrt{\frac{\gamma RT}{M}} = \sqrt{\frac{(1.40)(8.314\text{J/(molK)})(293\text{K})}{28.96 \times 10^{-3}\text{kg/mol}}} \approx 343\text{m/s}$$

Engineering check: $v \approx 331\text{m/s} + 0.6(20\text{m/s}) = 331 + 12 = 343\text{m/s}$. The two answers agree to the nearest metre per second.

Making sense of the answer: *The figure* 343m/s *is the value engineers, musicians and air-traffic controllers in every city in Tanzania use when they don't know the temperature precisely. At room temperature it is correct to better than half a per cent. Notice the two completely different lines of reasoning, namely a fundamental formula from kinetic theory and a hand-fitted linear rule from data, converging on the same answer. When two unrelated paths agree, a working physicist can carry on with confidence.*

Think like a physicist: *Every gas-temperature problem becomes a one-line substitution once* $v = \sqrt{(\gamma RT/M)}$ *is in your head. The two boxes you must fill before computing are the absolute temperature (always in kelvin, never in celsius) and the molar mass (always in kg per mole, never in g per mole). A wrong unit on* M *is the single most popular error on this topic on every examination paper that touches it.*

BINDER Example 4

Find the speed of sound in helium at $20°\text{C}$. Helium is monatomic with $\gamma = 5/3$ and $M = 4.00\text{g/mol}$. Then explain, in two sentences, why breathing helium makes the human voice sound much higher than usual.

Solution

At $T = 293\text{K}$ $(20°\text{C})$:

$$v_{He} = \sqrt{\frac{(5/3)(8.314\text{J/(molK)})(293\text{K})}{4.00 \times 10^{-3}\text{kg/mol}}} \approx 1007\text{m/s}$$

That is roughly $1007/343 \approx 2.94$ times the speed of sound in air at the same temperature, or about three times. When the air inside a person's vocal tract is replaced by helium, the resonant frequencies of the tract scale with the speed of sound inside it: each resonance now sits at roughly three times its normal frequency. The vocal cords themselves vibrate at the same fundamental frequency as before (the fold tension has not changed), but the higher resonances of the tract amplify the harmonics far more than usual, shifting the perceived timbre drastically upwards. The voice does not change its pitch in the strict musical sense, because the fundamental is unchanged, but the listener hears it as much higher because the harmonic mix has been radically reshaped.

Making sense of the answer: *Notice what changed and what did not. The helium has* $M = 4$ *instead of* 28.96*, a factor of about seven smaller; that pulls* v *up by a factor of about* $\sqrt{7} \approx 2.65$*. Helium also has* $\gamma = 5/3$ *instead of* $7/5$ *(a factor of* $25/21 \approx 1.19$ *inside the square root, i.e.* $\sqrt{1.19} \approx 1.09$ *on the speed). Multiplying gives* $2.65 \times 1.09 \approx 2.89$*, which lands inside one per cent of the full calculation. Two competing factors, lighter molecules and stiffer adiabatic behaviour, each push the speed up, and helium wins on both.*

Think like a physicist: *Always separate what depends on the gas (its* γ *and* M*) from what depends on the conditions (temperature). The trick of comparing speeds in two different gases is to take their ratio under the square root and let identical factors cancel. Whenever a problem asks for a ratio rather than a number, take the ratio early.*

HOT Example 5

A naval observer is positioned due north of a heavy gun, on flat ground, on a morning when the temperature is $20°\text{C}$. The observer sees the flash of the gun and starts a stopwatch; the report of the gun reaches her ears 4.0s later. A steady wind is blowing from east to west at 48km/h. Treating the speed of light as effectively infinite over the relevant distance, find the distance from the gun to the observer.

Solution

Speed of sound in still air at this temperature:

Using the engineering rule $v \approx 331 + 0.6T_C$:

$$v_{still} = 331\text{m/s} + 0.6 \times 20\text{m/s} = 343\text{m/s}$$

Decompose the wind:

The wind is from east to west; the observer is due north of the gun, so the sound is travelling due north on its way to her ear. The wind's velocity vector is along east–west; its component along the sound's direction of travel (north) is zero. Convert the wind speed for completeness:

$$w = 48\text{km/h} = \frac{48000\text{m}}{3600\text{s}} = 13.3\text{m/s}$$

The component of the wind in the direction the sound is travelling is the wind's full magnitude multiplied by the cosine of the angle between the wind vector and the line gun-to-observer. That angle is $90°$, so the component is $13.3\text{m/s} \times \cos 90° = 0$.

A perpendicular wind carries no energy along the sound's path; it advects the sound sideways but does not speed it up or slow it down along the gun-to-observer line.

Effective speed of sound towards the observer

Since the wind contributes nothing along the line of travel, the effective speed of sound from gun to observer is just the still-air value:

$$v_{eff} = v_{still} + 0 = 343\text{m/s}$$

Distance from the gun:
Distance equals effective speed times time of flight:

$$d = v_{eff} \times t = (343\text{m/s})(4.0\text{s}) = 1372\text{m} \approx 1.4\text{km}$$

Making sense of the answer: *Many students panic the moment a wind appears in a sound problem and automatically add or subtract the full wind speed to the speed of sound. The correct rule is: only the component of the wind along the line joining source and observer modifies the effective speed of sound between them. A wind blowing perpendicular to the line of travel contributes nothing to the time of flight, however violent it may be. The wind in this problem is doing two things: it is bending the sound's wavefronts sideways and it would shift the apparent direction the sound seems to come from, but it is not changing the time of flight from gun to ear.*

Think like a physicist: *Vector components are not a piece of mechanics nostalgia tacked onto sound for nostalgia's sake. They are the rule that decides which part of any vector influences any scalar quantity defined along a line. A wind is a vector. Time of flight is a scalar along the source-to-observer line. The rule that connects them is the dot product, and the dot product of a perpendicular wind with the line of travel is zero. Resolving the wind into components, along and across, is the safe move for every Doppler-and-wind problem in the rest of the chapter as well.*

With the speed of sound now in hand, sound has acquired its first hard number. The next section asks what happens when a sound wave is confined to a pipe rather than free to spread through open air. The pipe imposes a boundary condition at each end; the boundary condition selects a discrete set of allowed frequencies; and out of that selection falls the whole physics of organ pipes, flutes, clarinets, marimbas, and the resonance-tube experiment that you perform in the laboratory. Chapter 7's promise about pipes is the next one to be paid.

RESONANCE IN PIPES (OPEN AND CLOSED)

Sound that travels through open air spreads, weakens, and dies away within seconds. Sound trapped inside a pipe behaves very differently. The two ends of the pipe send the wave back the way it came, and the wave returning meets the wave still arriving. Most of the time the two cancel each other within a few cycles and the pipe falls silent. At certain special frequencies, however, the two reinforce each other perfectly and the pipe sings out a loud sustained note, picking that one frequency out of every other frequency the source might be supplying. These special frequencies are called the **resonant frequencies** of the pipe, and the phenomenon is called **acoustic resonance**. An organ pipe in a church sings on its own particular notes and no others, a flute plays one fingered pitch and refuses to play its neighbour, and the resonance tube experiment in your school laboratory is loud at one column length and almost silent at every neighbouring length. The pipe, in every case, is acting as a frequency filter, and it filters by the very same boundary-condition reasoning that Chapter 7 used for stationary waves on a stretched string.

Chapter 7 left us with two equivalent forms for the stationary-wave equation, the form with the spatial factor $\sin kx$ and the form with the spatial factor $\cos kx$, and explained that the boundary chooses the form. The sine form has a node at the position labelled $x = 0$ and is selected by a clamped boundary such as a string fixed to a wall. The cosine form has an antinode at $x = 0$ and is selected by a free boundary, where the medium is unconstrained. We promised then that the same boundary-condition reasoning would deliver the harmonic spectrum of an air column. The promise comes due here. The new ingredient is that a pipe can be sealed at one end and open at the other, or open at both ends, and the two boundary situations are not the same. A closed end traps the air and forbids displacement at that end. An open end leaves the air free to move but forbids any excess pressure there, because the air just outside the pipe is sitting at atmospheric pressure and the wave must match it. Out of those two simple rules falls the entire physics of pipe modes.

The closed pipe: one end sealed, one end open

Imagine a length of brass tubing held shut at the left end by a flat lid and left open at the right end. A clarinettist blows across the open end and pushes a sound wave into the tube. The wave runs to the lid, bounces back, and meets the next wave coming in. Some pairs of incoming and reflected waves will cancel each other within a fraction of a second and produce no sound. Others will reinforce each other and sustain a loud note. We want to find which wavelengths qualify.

Set up a coordinate system. Place the origin at the closed end and let $x = L$ mark the open end. The standing wave inside the pipe must satisfy two boundary conditions, one at each end, and the conditions are not the same.

At the closed end the air has nowhere to go: the lid is solid, the molecules cannot pass through it, and the displacement of the air there is forced to zero at every instant. That is a node of displacement at $x = 0$. The sine spatial form is the one that puts a displacement node at $x = 0$ without any further work, so we take it as our starting point:

$$s(x,t) = 2A\sin(kx)\cos(\omega t)$$

At the open end the air is free to move; the displacement is unconstrained. The pressure, however, is constrained. Just outside the open end the air sits at atmospheric pressure, and the small pressure variation produced by the wave must drop to zero there. The earlier section of this chapter on the nature of sound established that pressure and displacement are ninety degrees out of phase in a sound wave, which means a pressure node coincides with a displacement antinode. So the open end is a displacement antinode.

The mathematical condition for an antinode of the sine spatial form at $x = L$ is that $\sin(kL)$ equals plus or minus one. This in turn means kL must equal an odd multiple of $\pi/2$:

$$kL = (2n-1)\frac{\pi}{2}, \qquad n = 1,2,3,\ldots$$

Translating from wave number to wavelength using $k = 2\pi/\lambda$, the allowed wavelengths are:

$$\mathbf{L = (2n-1)\frac{\lambda_n}{4}}, \qquad \mathbf{n = 1,2,3},\ldots$$

and the allowed frequencies, found from $f = v/\lambda$, are:

$$\mathbf{f_n = \frac{(2n-1)\,v}{4L}}, \qquad \mathbf{n = 1,2,3},\ldots$$

The lowest of these is the fundamental:

$$\mathbf{f_1 = \frac{v}{4L}}$$

for which exactly one quarter of a wavelength fits inside the pipe, with a node at the closed end and the nearest antinode at the open end. The next allowed mode has $\mathbf{f_2 = 3f_1}$ and is called the ***first*** *overtone* or the *third harmonic*, with three-quarters of a wavelength inside the pipe. Above it sit $\mathbf{f_3 = 5f_1}$ (second overtone, ***fifth*** *harmonic*), $\mathbf{f_4 = 7f_1}$ (***third*** *overtone*, ***seventh*** *harmonic*), and so on. Only the odd harmonics show up. The even harmonics $2f_1$, $4f_1$, $6f_1$, and so on, are absent from the spectrum of a closed pipe because they cannot satisfy both boundary conditions at the same time. An even harmonic would demand either two nodes at the two ends or two antinodes, but a closed pipe insists on one of each.

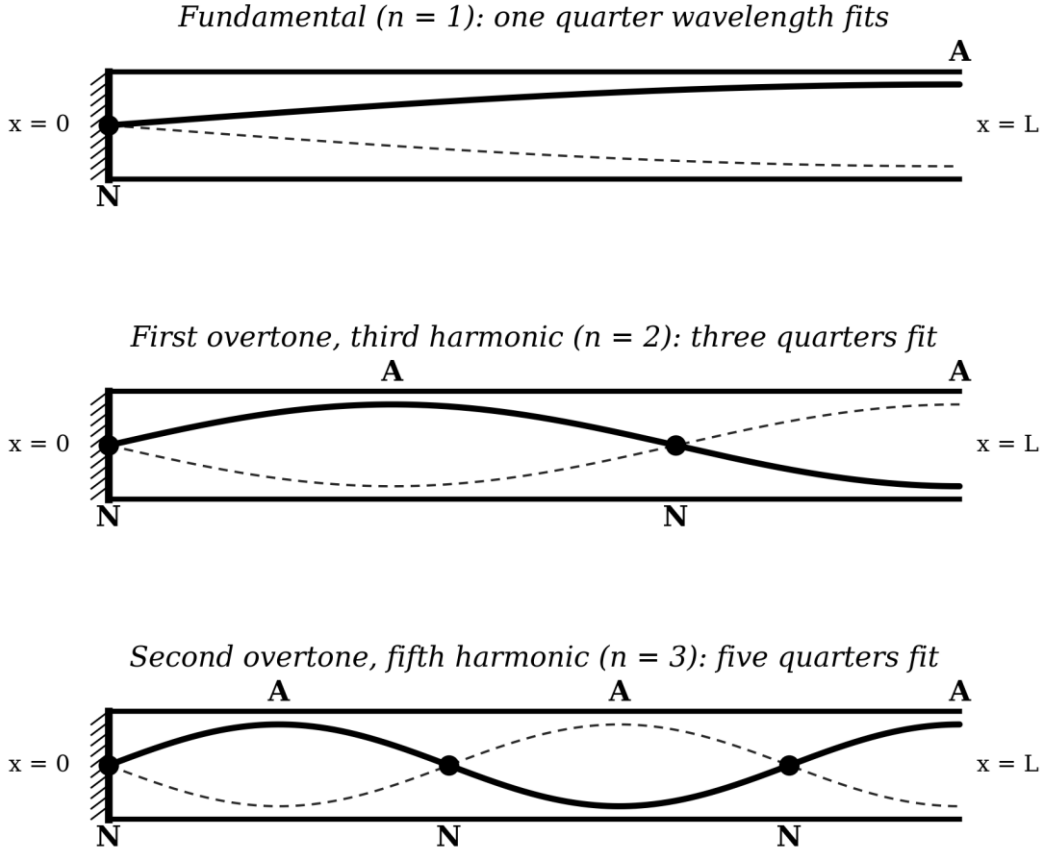


Figure: ***Standing-wave patterns of the first three modes of a closed pipe.*** *The closed end on the left is always a displacement node (N), and the open end on the right is always a displacement antinode (A).*

The open pipe: both ends open

Imagine now a length of brass tubing open at both ends, a piccolo or a flute or the barrel of a horn before anyone has stuck their hand into it. Place the origin $x = 0$ at one open end and let $x = L$ mark the other. Both ends are free; both ends are displacement antinodes; both ends are pressure nodes. The cosine spatial form is the one that puts a displacement antinode at $x = 0$ automatically, so it is the form we adopt:

$$s(x,t) = 2A\cos(kx)\cos(\omega t)$$

The condition for an antinode of the cosine spatial form at $x = L$ is that $\cos(kL)$ equals plus or minus one, which means kL must equal an integer multiple of π:

$$kL = n\pi, \qquad n = 1,2,3,\ldots$$

The allowed wavelengths are:

$$\mathbf{L = \frac{n\,\lambda_n}{2}, \qquad n = 1,2,3,\ldots}$$

and the allowed frequencies are:

$$\mathbf{f_n = \frac{n\,v}{2L}, \qquad n = 1,2,3,\ldots}$$

The fundamental of the open pipe is therefore $f_1 = v/(2L)$, exactly twice the fundamental of a closed pipe of the same length. At the fundamental, half a wavelength fits inside the open pipe, with antinodes at the two ends and a single node halfway between. The next allowed mode is $f_2 = 2f_1$, the first overtone or the second harmonic, with a complete wavelength inside the pipe. After that come $f_3 = 3f_1$, $f_4 = 4f_1$, and so on in unbroken succession. An open pipe is friendlier than a closed pipe, permitting all harmonics, both odd and even. This richer spectrum is the acoustic reason a flute (open at both ends) sounds brighter than a clarinet (closed at the reed end) on the same note: not louder, brighter, because the listener's ear is fed a fuller mix of higher partials, which the auditory cortex processes as timbre.

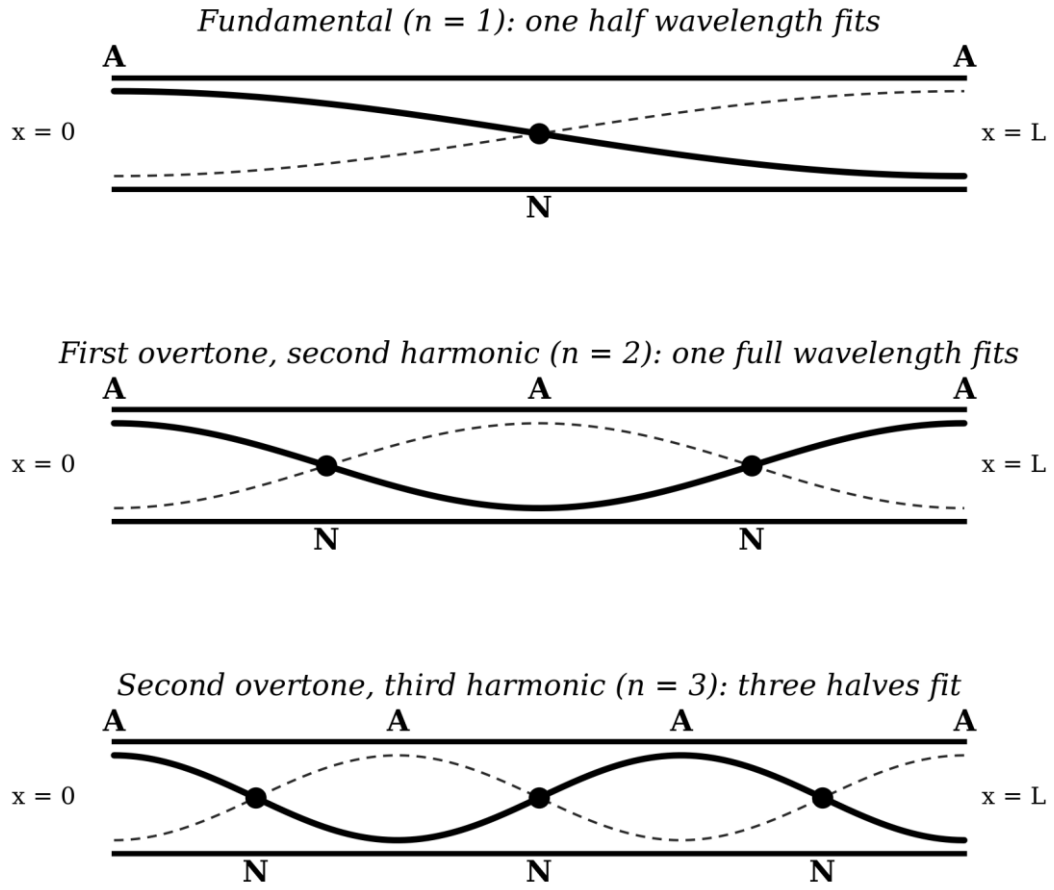


Figure: ***Standing-wave patterns of the first three modes of an open pipe.*** *Both ends are always displacement antinodes (A), with nodes (N) in between.*

Open against closed: the same length, two different worlds

Stand the two formulas side by side and the consequences are immediate. An open pipe and a closed pipe of the same physical length sing two different fundamentals:

$$f_1(\text{open}) = \frac{v}{2L}, \qquad f_1(\text{closed}) = \frac{v}{4L}$$

So the open pipe sounds one full octave above the closed pipe of equal length. Turning the comparison around, a closed pipe needs to be only half as long as an open pipe to sound the same fundamental, and that is one practical reason organ builders prefer closed pipes for the very lowest notes of a cathedral organ: a closed 16-foot pipe sounds at the same pitch as an open 32-foot pipe, at half the cost in lead and half the headroom needed in the loft above the choir. The price is the missing even harmonics, which give the closed-pipe stop a slightly hollow timbre next to the open-pipe stop. Different sounds for different musical purposes, but the same physics underwriting both.

The spacing of consecutive harmonics is different in the two cases as well. In an open pipe two consecutive harmonics differ by f_1, while in a closed pipe two consecutive harmonics differ by $2f_1$. A closed-pipe spectrum is therefore a sparse comb compared to an open-pipe spectrum, with only every other tooth present. That fact becomes the whole point of the HOT example a few paragraphs from now, in which the second overtone of a closed pipe is matched against the first overtone of an open pipe.

End correction: where the antinode really sits

All of the analysis so far has quietly assumed that the displacement antinode at the open end of a pipe lies exactly at the geometric end of the pipe. In a real pipe it does not. The air at the open end is free to move not only inside the pipe but also into the outside air just beyond the pipe, and so the antinode actually sits a small distance outside the geometric end. The correction is small but not zero, and any laboratory measurement of v using the resonance tube must account for it.

Experiment shows that the antinode of an open end of a pipe of radius r sits, to a good approximation, a distance

$$e \approx 0.6r$$

outside the geometric end. The effective acoustic length of the pipe is therefore longer than its physical length by e at each open end. For a closed pipe, with only one open end, the effective length is:

$$L_{\text{eff}} = L + 0.6r$$

and for an open pipe, with two open ends, the effective length is:

$$L_{\text{eff}} = L + 2(0.6r) = L + 1.2r$$

All the harmonic formulas written down above remain correct, provided L is replaced by the appropriate L_{eff} in each case. A clever way to bypass any uncertainty about e in a laboratory experiment is to determine v from two consecutive resonance lengths rather than one. The fundamental of a closed resonance tube satisfies $L_1 + e = \lambda/4$, and the very next resonance, the first overtone, satisfies $L_2 + e = 3\lambda/4$. Subtracting the first equation from the second makes e vanish entirely:

$$L_2 - L_1 = \frac{\lambda}{2}, \quad \text{so} \quad \boldsymbol{\lambda = 2(L_2 - L_1)}, \quad \text{and} \quad \boldsymbol{v = f\lambda = 2f(L_2 - L_1)}$$

Knowing the tuning-fork frequency f and the two resonance lengths L_1 and L_2, the experimenter computes v without ever having to know e. If the end correction is wanted afterwards, it is recovered by back-substituting into either resonance condition. This trick is what makes the resonance tube the canonical school-laboratory method for measuring the speed of sound in air.

Mistuned pipes: when two notes that should agree disagree

If two pipes that ought to sound the same note are sounded together but their lengths are slightly off, their fundamentals are slightly different and the listener hears beats. The beat frequency is the absolute difference of the two pipe frequencies, $f_b = |f_1 - f_2|$, exactly the formula Chapter 7 derived for two sinusoidal sources of nearly equal frequency. We use it here without re-deriving it and refer the reader who needs a refresher to the earlier chapter. A piano tuner counting beats against a tuning fork at Bagamoyo College of Arts is using the same arithmetic.

We now have five tools at our disposal: the closed-pipe spectrum, the open-pipe spectrum, the octave relationship between same-length pipes, the end-correction shortcut, and the beat formula for mistuned pipes. The worked examples that follow are five carefully chosen pieces of physics theatre designed to put each tool through its paces. Open your notebook; the pipes are about to start singing.

BINDER Example 6

A closed organ pipe is 0.68m long. Taking the speed of sound in air as 340m/s and neglecting end correction, find the wavelengths and frequencies of the three lowest modes of vibration.

Solution

For a closed pipe the spectrum is $f_n = (2n - 1)\, v/(4L)$, with $n = 1,2,3$, and so on. We work out each mode in turn.

Fundamental ($n = 1$):

$$\lambda_1 = 4L = 4 \times 0.68\text{m} = 2.72\text{m}$$
$$f_1 = \frac{v}{4L} = \frac{340\text{m/s}}{4 \times 0.68\text{m}} = 125\text{Hz}$$

First overtone ($n = 2$, third harmonic):

$$\lambda_2 = \frac{4L}{3} = \frac{2.72\text{m}}{3} = 0.907\text{m}$$
$$f_2 = 3f_1 = 3 \times 125\text{Hz} = 375\text{Hz}$$

Second overtone ($n = 3$, fifth harmonic):

$$\lambda_3 = \frac{4L}{5} = \frac{2.72\text{m}}{5} = 0.544\text{m}$$
$$f_3 = 5f_1 = 5 \times 125\text{Hz} = 625\text{Hz}$$

The three lowest modes have wavelengths of 2.72m, 0.907m, and 0.544m, and frequencies of 125Hz, 375Hz, and 625Hz.

Making sense of the answer: *The three frequencies sit in the ratio 1 : 3 : 5, the unmistakable fingerprint of a closed pipe. The pipe itself is only about the height of a school desk, yet it sounds as low as an open pipe more than 1.3m tall would. This is the closed-pipe bonus: by sealing one end, the builder doubles the effective acoustic length and saves half the brass.*

Think like a physicist: *When a closed-pipe problem turns up, write down the ratio 1 : 3 : 5 : 7 : … before you touch your calculator. Half of every closed-pipe question on an examination paper can be answered from the ratio alone, with no* $v/(4L)$ *ever computed. The 1 : 3 : 5 ratio is what distinguishes the closed-pipe spectrum from every other harmonic spectrum in physics, and recognising it on sight is worth a free mark in the right exam.*

BINDER Example 7

A closed organ pipe and an open organ pipe are sounded simultaneously, producing five beats per second. The open pipe is 30cm long. Find the change of length the closed pipe must undergo to bring the two into unison. Take the speed of sound in air as 340m/s and neglect end corrections.

Solution

Beats arise because the two fundamentals are close but not equal. The fundamental of the open pipe is fixed at:

$$f_o = \frac{v}{2L_o} = \frac{340\text{m/s}}{2 \times 0.30\text{m}} = 567\text{Hz}$$

The fundamental of the closed pipe must differ from this by 5Hz, so it is either 562Hz or 572Hz. The question does not say which, so we work out both possibilities.

Case 1: the closed pipe sounds 562Hz initially.

$$L_1 = \frac{v}{4f_1} = \frac{340\text{m/s}}{4 \times 562\text{Hz}} = 0.1512\text{m}$$

To bring the closed pipe to 567Hz, its length must be shortened to:

$$L_0 = \frac{v}{4 \times 567\text{Hz}} = \frac{340\text{m/s}}{4 \times 567\text{Hz}} = 0.1499\text{m}$$

The required change of length is:

$$\Delta L = L_1 - L_0 = 0.1512\text{m} - 0.1499\text{m} = 0.0013\text{m} \approx 1.3\text{mm shorter}$$

Case 2: the closed pipe sounds 572Hz initially.

$$L_1 = \frac{v}{4 \times 572\text{Hz}} = \frac{340\text{m/s}}{4 \times 572\text{Hz}} = 0.1486\text{m}$$

To bring the closed pipe to 567Hz, its length must be lengthened to $L_0 = 0.1499\text{m}$:

$$\Delta L = L_0 - L_1 = 0.1499\text{m} - 0.1486\text{m} = 0.0013\text{m} \approx 1.3\text{mm longer}$$

Either way the closed pipe needs adjustment by about 1.3mm.

Hence the change of length is 1.3mm.

Understand: Only the direction (lengthen or shorten) depends on whether the closed pipe was originally too sharp or too flat. To settle which case applies in practice, the player perturbs the closed pipe by a known small amount and listens to the beats. If the beat rate rises, the perturbation is moving away from unison and we are working the wrong way. If the beat rate falls, the perturbation is moving toward unison and we are working the right way.

Alternative solution

The two case-by-case computations both gave $|\Delta L| = 1.3\text{mm}$, hinting that the magnitude of the answer depends only on the magnitude of the beat frequency, not on the direction of the mistuning. This is exactly right and can be obtained in a single step.

For a pipe in any mode the fundamental frequency is inversely proportional to the length, $f = \text{constant}\left(\frac{1}{L}\right)$. Taking the differential of both sides and dividing through, the fractional change in frequency equals the fractional change in length, in magnitude:

$$\frac{|\Delta f|}{f} = \frac{|\Delta L|}{L}$$

The closed pipe, when brought to unison with the open pipe, has length

$$L = \frac{v}{4f} = \frac{340\text{m/s}}{4 \times 567\text{Hz}} = 0.1499\text{m}$$

And the required change of length therefore has magnitude:

$$|\Delta L| = L \times \frac{|\Delta f|}{f} = 0.1499\text{m} \times \frac{5\text{Hz}}{567\text{Hz}} = 1.32 \times 10^{-3}\text{m} \approx 1.3\text{mm}$$

The answer drops out in a single line, with no case branching. The branching is hidden in the absolute-value sign on Δf, which says that the same number of beats per second can arise from a closed pipe that is sharp or from one that is flat; the magnitude of the correction is the same, only its direction differs. To settle the direction in practice, perturb the closed pipe and listen, as in the case-by-case solution above.

Making sense of the answer: *Five hertz out of nearly six hundred is less than one per cent of the pitch, and 1.3mm out of nearly 150mm is also just under one per cent of the length. The two percentages agree because for a pipe in any mode the frequency is inversely proportional to the length, so* $|\Delta f/f|$ *equals* $|\Delta L/L|$*. The musician adjusting the tuning slide of a trombone is paying his rent on exactly this proportionality.*

Think like a physicist: *The absolute-value sign in the beat formula* $f_b = |f_1 - f_2|$ *hides important information: it does not tell the player whether his pipe is sharp or flat. To resolve the sign, perturb one of the two pipes in a known direction and listen to whether the beats speed up or slow down. This trick is the foundation of every piano-tuning method ever invented.*

REAL Example 8

A school marimba at Mwakaleli Secondary School in Mbeya is built from a row of closed tubes of varying length. One particular tube is designed to sound the note one octave above middle C, of frequency 523Hz.

(a) Compute the physical length of the tube required at a Mwakaleli rehearsal-room temperature of 25°C, neglecting end correction. Take the speed of sound at 25°C to be 346m/s.

(b) The marimba is tuned in Mwakaleli in January at 32°C, where the speed of sound is 350m/s. The same instrument is taken to Iringa in July, where the rehearsal hall sits at 15°C and the speed of sound is 340m/s. By what percentage does the pitch of the tube change, and is the resulting note sharp or flat? A physics teacher warns the class against the seductive but wrong reason.

Solution

(a) For a closed pipe at its fundamental, $L = v/(4f)$. At 25°C:

$$L = \frac{346\text{m/s}}{4 \times 523\text{Hz}} = 0.165\text{m} = 16.5\text{cm}$$

(b) Tuned at 32°C, the same physical length L gives a fundamental $f_m = 350/(4 \times 0.165) = 530\text{Hz}$. Carried to a cold Iringa hall at 15°C, with v dropping to 340m/s, the same length L gives:

$$f_i = \frac{340\text{m/s}}{4 \times 0.165\text{m}} = 515\text{Hz}$$

The pitch falls from 530Hz to 515Hz.

$$\text{Percentage drop} = \frac{f_m - f_i}{f_m} \times 100\% = \frac{530\text{Hz} - 515\text{Hz}}{530\text{Hz}} \times 100\% = 2.83\%$$

The percentage drop is 2.83%.

The note is **flat** in the cold Iringa hall.

Now the physics teacher intervenes.

Student: *"Sir, this makes perfect sense. In cold weather the brass tube contracts. A shorter tube gives a higher fundamental from* $f = v/(4L)$*, so the pitch should rise. But your calculation says the pitch falls. Have we made a sign error somewhere?"*

Physics teacher: *"Your reasoning is half right. The tube does contract in the cold. But before you cash in on that argument, work out how much. How much does a 16.5cm brass tube actually shorten between 32°C and 15°C?"*

Student: *"The linear expansivity of brass is about* 1.9×10^{-5} *per kelvin, from Chapter 2. The temperature drop is 17K. So* $\Delta L/L$ *is* $1.9 \times 10^{-5} \times 17$*, which is* 3.2×10^{-4}*, or about 0.03 per cent of the length."*

Mr. Physics teacher: *"Good. And the speed of sound? From the earlier section of this chapter it scales as the square root of absolute temperature. From 305K down to 288K, what does* v *do?"*

Student: *"*v *drops by the square root of the ratio.* $\sqrt{288/305}$ *is about 0.9717, so* v *falls by about 2.83 per cent."*

Physics teacher: *"Two effects are pushing on the pitch.* v *dropping pushes* f *down directly through* $f = v/(4L)$*.* L *contracting pushes* f *up through the same formula. The first effect weighs in at 2.83 per cent. The second effect weighs in at 0.03 per cent. The first effect is one hundred times bigger than the second. The marimba flattens in the Iringa hall because the speed of sound has dropped, not because the tube has shrunk. Always size up the dominant effect numerically before you let your intuition collect its prize."*

So the marimba arrives flat in Iringa by 2.83 per cent of its design pitch, with v doing almost all of the work. To restore the design pitch, the carpenter would need to trim each tube very slightly shorter to compensate; brass contraction by itself would not save the situation.

Making sense of the answer: *The same physics is the reason a wind instrument played in a cold lecture room sounds slightly flat after the air conditioning has been running, and the reason a band tunes up only after the players have warmed their instruments to room temperature. The lazy intuition that says "things shrink in cold weather, so the pitch goes up" is wrong for almost every wind and brass instrument made; the speed-of-sound effect dominates by orders of magnitude every single time.*

Think like a physicist: *Whenever two effects act on the same quantity, whether they are pulling in the same direction or, as here, in opposite directions, estimate each one numerically before concluding anything from intuition.*

REAL Example 9

A Form Five student at Kibaha Secondary School performs the resonance-tube experiment to measure the speed of sound in air. He sounds a tuning fork of frequency 480Hz above an air column whose length is varied by raising and lowering a water reservoir attached by a rubber hose. The first loud resonance is found when the air column is 17.2cm long. He raises the water level no further and instead lowers it to find the next resonance, which occurs at an air-column length of 53.4cm. The internal radius of the glass tube is 1.6cm.

(a) Compute the wavelength of the sound wave in the air column.

(b) Compute the speed of sound in the laboratory air on the day of the experiment.

(c) Compute the end correction, and compare it with the value $0.6r$ predicted from the radius of the tube.

Solution

(a) The first loud resonance of a closed pipe occurs when the air column is one quarter of a wavelength long (ignoring the end correction for a moment), and the next loud resonance occurs when the air column is three quarters of a wavelength long. The two true resonance conditions, including the end correction e at the open top of the tube, are:

$$L_1 + e = \frac{\lambda}{4}, \qquad L_2 + e = \frac{3\lambda}{4}$$

Subtracting the first equation from the second wipes the e out cleanly:

$$L_2 - L_1 = \frac{\lambda}{2}$$

So the wavelength is

$$\lambda = 2(L_2 - L_1) = 2 \times (53.4\text{cm} - 17.2\text{cm}) = 72.4\text{cm} = 0.724\text{m}$$

(b) The speed of sound follows from $v = f\lambda$:

$$v = f\lambda = 480\text{Hz} \times 0.724\text{m} = 347.5\text{m/s}$$

(c) Having found λ, the end correction is recovered from either of the two resonance conditions. Using the first:

$$e = \frac{\lambda}{4} - L_1 = \frac{0.724\text{m}}{4} - 0.172\text{m} = 0.009\text{m} = 0.9\text{cm}$$

The $0.6r$ rule predicts:

$$e_{\text{predicted}} = 0.6 \times r = 0.6 \times 1.6\text{cm} = 0.96\text{cm}$$

The experimental and predicted values agree to within about six per cent, well inside the precision of a school laboratory.

Making sense of the answer: *The speed of sound came out at 347.5m/s, comfortably inside the band a Tanzanian classroom should expect, where a temperature around 27°C gives a sound speed close to 347m/s. The measurement would not have been possible at all without the end-correction trick: if the student had used only the first resonance and assumed $L_1 = \lambda/4$, he would have computed $\lambda = 0.688m$, $v = 330.2m/s$, and would have been low by about five per cent. Subtracting two resonance lengths to eliminate e buys back exactly that error.*

Think like a physicist: *Whenever an experiment contains a small unknown systematic offset, look for a way to cancel it by taking two measurements and subtracting one from the other. The two-resonance-lengths trick here is a special case of a much more general experimental technique: never measure an absolute quantity if a difference of two absolute quantities will do the same job. A Form Five physics student who internalises this rule once will use it for the rest of his life in every laboratory he ever sets foot in..*

HOT Example 10

A closed organ pipe and an open organ pipe are arranged so that the second overtone of the closed pipe sounds at the same frequency as the first overtone of the open pipe. Find the ratio of the length of the closed pipe to the length of the open pipe. Given that the common frequency is 880Hz and that the speed of sound in air is 340m/s, find each length. Neglect end corrections.

Solution

Write each frequency in turn. The closed-pipe spectrum is $f_n = (2n - 1)\, v/(4L_c)$, with $n = 3$ for the second overtone, which is the fifth harmonic:

$$f_{\text{2nd overtone, closed}} = \frac{5\,v}{4L_c}$$

The open-pipe spectrum is $f_n = n\, v/(2L_o)$, with $n = 2$ for the first overtone, which is the second harmonic:

$$f_{\text{1st overtone, open}} = \frac{2\,v}{2L_o} = \frac{v}{L_o}$$

Setting the two equal:

$$\frac{5\,v}{4L_c} = \frac{v}{L_o}$$
$$\frac{L_c}{L_o} = \frac{5}{4}$$

The closed pipe is therefore five-quarters as long as the open pipe. With a common frequency of 880Hz, the open pipe satisfies:

$$L_o = \frac{v}{f} = \frac{340\text{m/s}}{880\text{Hz}} = 0.386\text{m} = 38.6\text{cm}$$

and the closed pipe satisfies:

$$L_c = \frac{5}{4} L_o = 1.25 \times 0.386\text{m} = 0.483\text{m} = 48.3\text{cm}$$

The two pipes have lengths 38.6cm (open) and 48.3cm (closed).

Making sense of the answer: *Half the trick in this question lies in matching the right overtone of each pipe to the right multiple of the fundamental. The second overtone of a closed pipe is the fifth harmonic, not the third, because the closed-pipe spectrum starts at the first harmonic and then skips every even multiple. The first overtone of an open pipe is the second harmonic, because the open-pipe spectrum starts at the first harmonic and then proceeds through every consecutive integer multiple. Mixing up the overtone-versus-harmonic vocabulary is the most common error on this kind of question, and the examiner who set it knew exactly that.*

Think like a physicist: *When a pipe problem mixes a closed pipe and an open pipe, write the two spectra side by side before substituting a single number: closed pipe $f_n = (2n - 1)\, f_1(closed)$ for $n = 1,2,3,$ ..., and open pipe $f_n = n\, f_1(open)$ for $n = 1,2,3,$ .... The substitutions then become mechanical and almost no algebra is required. The physics is hidden in the boundary conditions the two spectra encode, not in the arithmetic that follows.*

We now have a fairly complete theory of sound. We can compute its speed at any temperature in any gas, predict the harmonic spectrum of any pipe of any length, account for the small end correction at every open end, and read off the beat frequency between any two mistuned pipes that share a stage. But we have not yet asked one question, a question the human ear answers automatically every waking second: *how loud is loud?* The next section introduces the logarithmic scale that engineers, physiologists, and examiners all use for that purpose.

LOUDNESS, INTENSITY, AND THE DECIBEL SCALE

The earlier section of this chapter on the nature of sound listed three things the ear hears: pitch, loudness, and quality. It also paid two of those debts immediately. Pitch was identified with frequency. Quality was identified with the harmonic spectrum. Loudness was given the word "intensity" and then quietly postponed, on the grounds that intensity has its own mathematics and that the mathematics deserves a section of its own. The promise comes due now. This section answers, once and for all, the question the ear answers every waking second without thinking: how loud is loud?

The question turns out to be harder than it looks. The human ear can detect a sound as soft as a mosquito flying three metres away in a quiet room and can tolerate, with discomfort but without immediate injury, a sound as loud as a jackhammer one metre away. Measured in watts per square metre, the gap between these two sounds is a factor of a million million. Writing the loud sound as one million million times the soft sound is unwieldy on a page. Asking a Form Five student to feel comfortable with such a number is unreasonable. And the human nervous system itself, conveniently, does not perceive sound on a linear scale; it perceives sound approximately logarithmically, so that successive equal doublings of intensity feel like successive equal steps of loudness.

Engineers, physiologists, and the writers of every reasonable physics syllabus in the world have therefore agreed on a single logarithmic scale for sound: the **decibel**. This section introduces it, derives the two pieces of arithmetic the student must drill until they need no thought, applies it to a point source whose listener walks away from it, and applies it to a roomful of incoherent sources whose intensities add even though their levels do not. Five or six pages from now you will hold the entire decibel toolkit.

Intensity: the physical quantity behind the perception

Chapter 7 introduced the intensity of any wave as the power the wave carries across a unit area perpendicular to its direction of travel, measured in watts per square metre. The same chapter showed two further things that we will use heavily here and will not re-derive. First, the intensity of a sinusoidal wave is proportional to the square of its amplitude, $I \propto A^2$. Second, the intensity of a wave radiating outwards from a small source in all directions falls off as the inverse square of the distance from the source,

$$I = \frac{P}{4\pi r^2}$$

where P is the total power radiated and r is the distance from the source to the listener. Both results are taken as given. The reader who needs a refresher should glance back at Chapter 7 before proceeding; everything in this section sits on top of those two facts.

The softest sound a healthy young human ear can detect at the frequency 1kHz is an intensity of about

$$\mathbf{I_0 = 10^{-12}\ W/m^2}$$

This number is called the **threshold of hearing** and is the reference point against which every other sound intensity is to be compared. It is not a fundamental constant of physics; it is a fact about the human ear, measured by carefully controlled experiments on hearing-test volunteers. A different species, a younger ear, a different frequency, or a slightly damaged eardrum would give a slightly different threshold. But the convention of physics is fixed, and the convention is the one above.

The loudest sound the human body tolerates briefly, without immediate injury, is about

$$\mathbf{I_{pain} \approx 1\ W/m^2}$$

This number is called the **threshold of pain**. Above it, exposure for even a few seconds risks permanent hearing loss. Between the threshold of hearing and the threshold of pain the ear must work over a span of

$$\frac{I_{pain}}{I_0} = \frac{1}{10^{-12}} = 10^{12},$$

twelve full orders of magnitude. No linear scale can handle that span on a single page without absurdity. The next subsection explains how the logarithmic scale handles it gracefully.

Why a logarithm? The twelve-decade problem

If a single number must describe sounds that vary by a factor of 10^{12}, the number had better be a logarithm. Logarithms convert multiplication into addition and convert wide ratios into narrow numbers. The base-10 logarithm of 10^{12} is exactly 12, a simple value that is easy to handle. The base-10 logarithm of 2 is about 0.301, and the base-10 logarithm of 10 is exactly 1. These three facts, $\log 2 \approx 0.301$, $\log 10 = 1$, and the general identity $\log(ab) = \log a + \log b$, between them, are all the logarithm arithmetic this section will use.

The logarithmic measure of sound intensity goes by the name of Alexander Graham Bell, the inventor of the telephone, who first proposed it for use in long-distance telephone engineering. A factor of ten in intensity is called one bel. The bel itself turns out to be a slightly inconvenient unit, because the practical range of human hearing is only twelve bels wide. A unit one tenth as large gives a more comfortable working range, from 0 to 120, and the unit one tenth as large is called the **decibel**, abbreviated **dB**. The decibel is the working unit of loudness throughout the physics, engineering, and physiology of sound.

The decibel formula

The intensity level L of a sound of intensity I is defined by

$$\mathbf{L = 10 \log_{10}\left(\frac{I}{I_0}\right)} \quad \textit{decibels}$$

where $I_0 = 10^{-12}\ \text{W/m}^2$ is the threshold of hearing introduced above. Several things about this definition deserve immediate comment.

First, *L is dimensionless.* The argument of the logarithm is a ratio of two intensities and so is itself a pure number. The factor of 10 in front is a pure number. The unit "decibel" is an honorary tag, not a physical dimension, and a decibel value cannot be converted to a number of joules or watts the way a kilogram can be converted to a number of newtons. A decibel is a label, not a quantity, in the dimensional sense.

Second, *the threshold of hearing sits at $L = 0$ exactly,* because $\log(I_0/I_0) = \log 1 = 0$. A sound at the threshold of hearing is, by definition, at 0 decibels. A quieter sound has a negative decibel level, which is a perfectly valid number, though for the human ear at 1kHz it represents a sound that nobody can hear.

Third, *the threshold of pain sits near $L = 120$,* because $\log(1/10^{-12}) = \log 10^{12} = 12$, and $10 \times 12 = 120$. A sound at the threshold of pain is about a million million (one trillion) times more intense than a sound at the threshold of hearing, and the dB scale lets a Form Five student write this comparison as "120dB" without flinching.

Fourth, *the dB scale is not an absolute measure of loudness;* it is a measure of intensity relative to the threshold of hearing. A different reference would give different numerical levels for the same physical sounds. In acoustics, in electrical

engineering, in optics, and in radio engineering, several other references are in use for different purposes, all of them written with the same formula but with a different I_0. The reference for sound in air is the one above, and is the only reference this chapter uses.

Two patterns to drill into your fingers

Two arithmetic patterns recur in dB problems often enough that the practised student does them mentally, with no calculator, and the unpractised student loses minutes to them on every examination.

The first pattern is doubling. When the intensity doubles, $I \to 2I$, the dB level increases by

$$\Delta L = 10\log\left(\frac{2I}{I}\right) = 10\log 2 \approx 10 \times 0.301 = 3.01\text{dB} \approx 3\text{dB}$$

So doubling intensity adds three decibels, and halving intensity subtracts three decibels. This is the most useful single fact in dB arithmetic. A practising audio engineer in a recording studio at Mlimani City, switching from one speaker to two identical speakers, knows in advance that the listening room will get exactly 3dB louder, no measurement required.

The second pattern is ten-fold. When the intensity multiplies by ten, $I \to 10I$, the dB level increases by

$$\Delta L = 10\log\left(\frac{10I}{I}\right) = 10\log 10 = 10 \times 1 = 10\text{dB}$$

So a ten-fold intensity adds ten decibels, and a hundred-fold intensity adds twenty decibels, and a thousand-fold intensity adds thirty decibels, and so on, with one extra factor of 10 in intensity worth ten extra dB on the level scale every time.

Combining the two patterns lets the student handle every common intensity ratio without using a calculator. An intensity ratio of 4 is two doublings, worth $2 \times 3 = 6\text{dB}$. An intensity ratio of 8 is three doublings, worth $3 \times 3 = 9\text{dB}$. An intensity ratio of 20 is a ten-fold (10dB) followed by a doubling (3dB), worth 13dB. An intensity ratio of 100 is two ten-folds, worth 20dB. An intensity ratio of 1000 is three ten-folds, worth 30dB. The student who has these conversions in his fingers will read most dB problems straight off the page.

Drop with distance: the six-decibel rule

A point source radiates power P in all directions equally. From the inverse-square law,

$$I = \frac{P}{4\pi r^2}$$

The intensity at distance r falls off as $1/r^2$. The dB level at distance r relative to the level at some reference distance r_1 is therefore:

$$L_2 - L_1 = 10\log\left(\frac{I_2}{I_1}\right) = 10\log\left(\frac{r_1^2}{r_2^2}\right) = 20\log\left(\frac{r_1}{r_2}\right)$$

If $r_2 > r_1$ then $r_1/r_2 < 1$, the logarithm is negative, and the level drops as expected. If $r_2 < r_1$ the listener has moved closer and the level rises. Either way, the dB change with distance is given by:

$$\Delta L = L_2 - L_1 = 20\log\left(\frac{r_1}{r_2}\right) = -20\log\left(\frac{r_2}{r_1}\right),$$

with the sign of the change determined by which way the listener walks.

The single most useful consequence of this formula is the six-decibel rule for point sources. Doubling the distance from a point source, $r_2 = 2r_1$, drops the level by

$$|\Delta L| = 20\log 2 \approx 20 \times 0.301 = 6.02\text{dB} \approx 6\text{dB}$$

A pedestrian walking along a quiet street, moving from 10 metres away from a loudspeaker to 20 metres away, will experience a drop in sound level of about six decibels, even though the source of the sound remains unchanged.

The six-decibel rule applies only to idealised point sources in open air with no reflections. A real urban street is full of reflections from buildings on either side, and the rule overestimates how quickly the sound dies with distance. A real concert hall is designed to keep reflections alive so that the audience at the back of the hall hears almost the same loudness as the audience near the front. Reverberation and reflection are a topic for an applied acoustics chapter that this book will not write; for our purposes, the six-decibel rule is a good first approximation and an excellent examination tool.

Combining incoherent sources: intensities add, levels do not

If two sound sources are sounded together and are incoherent, meaning that their phases bear no fixed relationship to each other, their intensities add. This is one of the basic results of Chapter 7 on interference and superposition. We use it here without re-deriving it.

For two incoherent sources of intensities I_1 and I_2 at the listener's ear, the total intensity is:

$$I_{\text{total}} = I_1 + I_2$$

The total level L_{total} is then computed from this total intensity through the usual formula. The levels themselves do not add. A common student error is to say, of two simultaneous 50dB sources, that the combined level is 100dB. The combined level is in fact 53dB. Here is why: two equal-intensity sources have a combined intensity of 2I, which is a doubling, which adds 3dB to the level. Two equal sources at 50dB combine to a single level of 53dB.

More generally, n equal-intensity sources combine to a total intensity of nI and a total level higher by:

$$\Delta L = 10\log n$$

decibels. So three identical sources add $10\log 3 \approx 4.8$ dB. Ten identical sources add exactly 10dB. A hundred identical sources add exactly 20dB. The rule scales smoothly with n and applies in every situation where the sources are incoherent, which in practice is every situation outside of carefully arranged laboratory interference experiments.

Equal-loudness contours and the dBA tag

A complication the dB scale does not capture is that the ear is not equally sensitive to all frequencies. At the same physical intensity, a 1kHz tone sounds substantially louder than a 30Hz tone, because the ear is most sensitive in the speech-frequency range from about 1 to 4kHz and rolls off in sensitivity at both lower and higher frequencies. The standard plot of equal perceived loudness versus frequency and intensity, called the **equal-loudness contour**, looks like a stack of U-shaped curves, with each curve carrying its own loudness label. The Form Five student is not required to draw or work with these curves but should know they exist.

To handle this frequency dependence approximately in noise-control engineering, a frequency-weighted version of the decibel scale is used. It is called the A-weighted decibel and is denoted dBA. The dBA value of a sound is its physical dB value, adjusted by the standardised A-weighting filter, so that a 60dBA reading corresponds to a sound that a typical listener would perceive as if it were 60dB at 1kHz. Industrial hearing-protection rules, city noise ordinances, and most legal noise limits are written in dBA, not in dB. A typical Tanzanian municipal council writes residential noise limits as "60dBA during the day, 45dBA at night." For the purposes of this book, the dBA is mentioned only so that students are not surprised by it when they encounter it.

A reference table for the loud and the soft

The table below anchors the dB scale to sounds the student will recognise. Memorising the table is not necessary; recognising at a glance whether a quoted level is realistic is.

Sound	Approximate intensity (W/m²)	Level (dB)
Threshold of hearing (1kHz, healthy young ear)	10^{-12}	0
Rustle of leaves at 5m	10^{-10}	20
Quiet conversation at 1m	10^{-7}	50
Daladala at full speed, passenger view	10^{-4}	80
Construction site at 5m	10^{-2}	100
Threshold of pain	1	120
Saba-Saba exhibition loudspeaker stack at 1m	10^{2}	140 (instant injury)

A passenger on a daladala in afternoon traffic on Morogoro Road is sitting in a sound field of about 80dB. A conversation across the aisle takes place against a background that is one hundred million times more intense than what the same passenger would call quiet, yet by the logarithmic standard of the dB scale it is "only" 80dB. The scale's compression is what makes it usable at all.

We now have five pieces of equipment in the dB toolkit. The worked examples that follow are three carefully chosen pieces of physics theatre designed to exercise each of those pieces in turn, with the third example landing on a surprise that catches more students by misjudgement than any of the previous five worked examples in this chapter combined.

BINDER Example 11

The intensity of a sound at a particular point in a Mwenge auditorium is 5.0×10^{-5} W/m^2. Find the intensity level in decibels. Take the threshold of hearing as $I_0 = 10^{-12}$ W/m^2.

Solution

Apply the decibel formula directly:

$$L = 10\log\left(\frac{I}{I_0}\right) = 10\log\left(\frac{5.0 \times 10^{-5}\ \text{W/m}^2}{10^{-12}\ \text{W/m}^2}\right) = 76.99\text{dB} \approx 77\text{dB}$$

The intensity level is therefore about 77dB, in the same neighbourhood as a daladala or a noisy classroom.

Making sense of the answer: *The student who has internalised the two arithmetic patterns can verify this answer without ever using a calculator. The intensity ratio* 5.0×10^7 *is half of* 10^8*, which would be 80dB exactly. Half of an intensity is a loss of 3dB. So* 5.0×10^7 *is* $80 - 3 = 77$ *dB. Two patterns and one subtraction, no logarithm key pressed.*

Think like a physicist: *Always estimate before computing. A 77dB answer is reasonable for the inside of a busy room; a 7dB or a 770dB answer would be obviously wrong. The simple check "where does this answer lie between 0 and 120?" catches almost every such error before any marks are lost.*

REAL Example 12

A road construction site near Magomeni produces a sound level of 95dB at a distance of 10m from its centre. Tanzanian occupational safety guidelines require workers to wear hearing protection at any location where the sound level reaches or exceeds 85dB.

(a) Find the minimum distance from the centre of the site at which a worker without hearing protection just complies with the rule, assuming the site behaves as a point source in open air.

(b) Show that doubling the distance from a point source drops the level by approximately 6dB.

Solution

(a) The dB level at distance r_2 relative to the level at a reference distance r_1 follows the inverse-square law applied to dB. The change in level is

$$\Delta L = L_2 - L_1 = -20\log\left(\frac{r_2}{r_1}\right)$$

Substitute $L_1 = 95\,\text{dB}$ at $r_1 = 10\,\text{m}$, and require $L_2 = 85\,\text{dB}$ at the unknown r_2:

$$\Delta L = 85\text{dB} - 95\text{dB} = -10\text{dB}$$

Hence:

$$-10\,\text{dB} = -20\log\left(\frac{r_2}{10\,\text{m}}\right)$$

$$r_2 = 31.62\,\text{m}$$

A worker without hearing protection must stand at least 31.6m from the centre of the site to comply with the rule.

(b) Doubling the distance means $r_2 = 2r_1$. Substitute into the formula:

$$\Delta L = -20\log(2) = -20 \times 0.301 = -6.02\,\text{dB} \approx -6\,\text{dB}$$

Hence doubling the distance drops the level by approximately 6 decibels. (This is the six-decibel rule for point sources.)

Making sense of the answer: *Bringing the level down by 10dB requires the ratio of intensities to drop by a factor of 10, which by the inverse-square law requires the distance to grow by a factor of* $\sqrt{10} \approx 3.16$. *The factor of* $\sqrt{10}$ *arises because intensity goes as* $1/r^2$*, so a ten-fold reduction in intensity is a* $\sqrt{10}$*-fold growth in distance. The numerical answer of 31.6m is reasonable for a road construction site; it agrees with the rule of thumb that safety engineers in noisy industries use mentally without recourse to a calculator.*

Think like a physicist: *The six-decibel rule is not a coincidence of arithmetic; it is the inverse-square law dressed up in logarithmic clothing. Whenever a question about a point source mentions doubling, halving, or a factor of any power of two in either distance or intensity, the answer will involve an integer multiple of 3dB (for intensities) or an integer multiple of 6dB (for distances). Reaching for these mental shortcuts before reaching for the calculator is the difference between a student who finishes the dB question with time to spare and a student who runs out of paper at the end.*

HOT Example 13

A primary-school classroom at Mawenzi Primary School in Moshi has an average background sound level of 35dB when the children are quiet. With 36 children present and chattering, the average level rises to 65dB. Treating each child as an independent, incoherent source of equal average intensity, and ignoring the small contribution of the background sound to the chattering total,

(a) estimate the average sound intensity contributed by one child, in W/m², averaged over the room;

(b) predict the sound level the teacher would measure if the class size doubled to 72 children, with each child contributing the same intensity as before.

Solution

(a) The two given dB levels must be converted back to intensities. The decibel formula $L = 10\log(I/I_0)$ can be solved for I by taking the inverse logarithm of both sides.

$$L = 10\log\left(\frac{I}{I_0}\right)$$

$$\mathbf{I = I_0 \log^{-1}\left(\frac{L}{10}\right)}$$

This is the inverse-log form of the decibel formula and is the form a student must reach for whenever a dB level is given and an intensity is wanted. Apply it to the background sound at $L = 35$ dB:

$$I_{\text{background}} = 10^{-12}\,\text{W/m}^2 \log^{-1}\left(\frac{35}{10}\right) = 3.16 \times 10^{-9}\,\text{W/m}^2$$

Apply it again to the total sound at $L = 65\,\text{dB}$:

$$I_{\text{total}} = 10^{-12}\,\text{W/m}^2 \log^{-1}\left(\frac{65}{10}\right) = 3.16 \times 10^{-6}\,\text{W/m}^2$$

The intensity coming from the children alone is the total intensity minus the background intensity:

$$I_{\text{children}} = I_{\text{total}} - I_{\text{background}} = 3.16 \times 10^{-6}\,\text{W/m}^2 - 3.16 \times 10^{-9}\,\text{W/m}^2$$

The background contributes one part in a thousand to the total and can safely be neglected:

$$I_{\text{children}} \approx 3.16 \times 10^{-6}\,\text{W/m}^2$$

Divide by the number of children:

$$I_{\text{per child}} = \frac{I_{\text{children}}}{36} = \frac{3.16 \times 10^{-6}\,\text{W/m}^2}{36} = 8.78 \times 10^{-8}\,\text{W/m}^2$$

Each child contributes, on average, about $8.78 \times 10^{-8}\,\text{W/m}^2$ of intensity to the room.

(b) With 72 (2×36) children, each contributing the same average intensity as before, the combined intensity from the children alone doubles:

$$I_{\text{children, new}} = 2 \times I_{\text{children}} = 2 \times 3.16 \times 10^{-6}\,\text{W/m}^2 = 6.32 \times 10^{-6}\,\text{W/m}^2$$

Now go back through the forward decibel formula to find the corresponding level. Ignoring the negligible background,

$$L_{\text{total, new}} = 10\log\left(\frac{I_{\text{children, new}}}{I_0}\right) = 10\log\left(\frac{6.32 \times 10^{-6}\,\text{W/m}^2}{10^{-12}\,\text{W/m}^2}\right) = 68.01\text{dB} \approx 68\text{dB}$$

Doubling the class size raises the level by only 3dB, from 65dB to 68dB.

Mr. Akilikubwa walks in at this moment and asks the class, with a face he uses when he expects somebody to be surprised.

Mr. Akilikubwa: *"Kipanga, are you content with this answer?"*

Kipanga: *"Sir, I am not. We have doubled the children. Why has the level risen by only 3dB? Surely twice as many children should be twice as loud."*

Mr. Akilikubwa: *"And what does 'twice as loud' mean to the ear, Kipanga? On the dB scale, twice the intensity adds 3dB. We doubled the intensity, so we added 3dB. The ear, on the other hand, does not perceive 3dB as a doubling of loudness; it perceives a doubling of loudness as roughly 10dB, which requires not a doubling of children but a tenfold multiplication of children. To make the classroom sound twice as loud to the teacher, you would need not 72 children but about 360 of them."*

Kipute: *"So the head teacher who complains that the school is twice as noisy as last year is probably wrong by a factor of about five."*

Mr. Akilikubwa: *"She is probably wrong by a factor of about five, yes, but only because she is using her ear, which is logarithmic, while her impression speaks in linear language. The dB scale is closer to what the ear actually feels; the linear count of children is closer to what the eye actually counts. The two scales disagree by exactly the amount the logarithm compresses."*

Making sense of the answer: *The answer follows the doubling-adds-3-dB rule exactly, applied not to a single source but to a collection of equal incoherent sources. The general result is that n equal sources combine to a level $10\log n$ dB above the level of one such source, and so 72 children sit $10\log(72/36) = 10\log 2 = 3.01$ dB above 36 children, regardless of what the individual per-child intensity is. The answer would have been the same if each child had been twice as loud or half as loud, because the doubling-rule depends only on the ratio of the numbers of sources, not on the absolute intensities. This is one of the cleanest demonstrations in the chapter of why the dB scale is the right scale: it makes the physically relevant ratios visible without forcing the student to track unwieldy absolute numbers.*

Think like a physicist: *The classroom example is the simplest case of a result that scales to large auditoria, traffic streams, factories, and music festivals. Whenever the engineer wants to halve the perceived noise of a workplace, halving the number of noise sources will not do; the engineer must reduce the number of sources by a factor of about ten, or alternatively reduce the intensity of each source by the same factor. This is why noise-control engineering in industry focuses overwhelmingly on quieter individual machines rather than on fewer machines: linear cuts in count do not pay for themselves in dB savings, but logarithmic cuts in intensity per machine do. Once you see the world in dB you cannot un-see it, and the world is full of acoustic surprises waiting for somebody to notice them.*

We have now answered, in physical terms, the question the ear answers every waking second without thinking. We can compute the intensity level of any sound from its intensity, predict how its level falls off as the listener walks away from the source, and combine the levels of many incoherent sources without confusing dB with intensity. The next section turns away from sound altogether and goes to the other great family of waves the chapter must cover: the electromagnetic family, in which the speed is not 340m/s but 3×10^8m/s, the medium is not air but the vacuum of empty space, and the wave variable is not pressure but electric and magnetic field. Sound has been our friend for four sections; it is time to meet its faster, stranger cousin.

ELECTROMAGNETIC WAVES (EM) AND THE SPEED OF LIGHT (c)

The first four sections of this chapter belonged to sound. We described its nature, found its speed in air, identified the resonant frequencies of any pipe of any length, and converted intensity into the logarithmic units the human ear actually uses. Sound, however, is not the only wave the syllabus knows. The wave that lets the reader see the page in front of him, the wave that carries the voice of a TBC Taifa newsreader from Dar es Salaam to a receiver in Musoma, the wave that lets a doctor at Muhimbili photograph a patient's broken bone, and the wave that defines the upper speed limit of the universe, all belong to a single family. The family is called the **electromagnetic** family. Unlike sound, its members do not need a medium to propagate; they travel through the vacuum of empty space at a fixed speed, the same speed in every direction, the same speed for every observer, and the same speed for every member of the family from radio waves to gamma rays. The chapter that gave four sections to sound must now give equal billing to its faster, stranger cousin.

In this section we do six things. We state Maxwell's great prediction of 1865, namely that electric and magnetic fields can propagate through empty space as a coupled wave at a speed determined entirely by the two constants μ_0 and ϵ_0. We compute that speed. We list the properties of the resulting wave, which the student will find both familiar (the wave equation is the same form Chapter 7 derived for strings) and strange (no medium is required, and the wave behaves identically in glass and in water with only the constants adjusted). We write down the sine forms of E and B that show how the wave actually propagates through space and time. We note how the speed adjusts when the wave passes from vacuum into a material. We work through the experimental measurement of c performed by Léon Foucault in 1862, the cleanest school-laboratory method ever devised. By the end of the section the student will be able to compute c from μ_0 and ϵ_0, identify the structure of a plane electromagnetic wave, write its sine-form expressions in time and space, compute its intensity from the amplitude of its electric field, find the speed of light in a transparent material from its refractive index, and read a Foucault apparatus question off the page.

Maxwell's prediction: light is an electromagnetic wave

In 1865 the Scottish physicist James Clerk Maxwell published a paper called *A Dynamical Theory of the Electromagnetic Field*. In it he showed that the four equations governing electric and magnetic fields, today known as Maxwell's equations (Gauss's law for electricity, Gauss's law for magnetism, Faraday's law of induction, and Ampère's law with Maxwell's own correction term), admit wave solutions in vacuum. The full derivation of these wave solutions requires Maxwell's equations in differential form, which lies beyond the scope of this book. Instead we state the result and identify, by direct comparison with a wave equation the student already knows, the speed of the resulting wave.

We begin with the wave equation the student already knows. Chapter 7 showed that a small transverse displacement $y(x,t)$ on a stretched string obeys:

$$\frac{\partial^2 y}{\partial t^2} = v^2 \frac{\partial^2 y}{\partial x^2}$$

The coefficient sitting in front of the spatial second derivative on the right-hand side is the squared wave speed.
Maxwell's equations, when manipulated in a region of empty space where no charges and no currents are present, produce a wave equation of exactly this same form for the electric field. The result is:

$$\frac{\partial^2 E}{\partial t^2} = \frac{1}{\mu_0 \epsilon_0} \frac{\partial^2 E}{\partial x^2},$$

with an identical equation holding for the magnetic field. Comparing this electromagnetic wave equation with the string wave equation term by term, the coefficient sitting in front of the spatial second derivative on the right-hand side is now $1/(\mu_0 \epsilon_0)$ in place of v^2. The two equations have the same form, so the role of the coefficient is the same in both: the coefficient is the squared wave speed of whatever wave the equation describes. We therefore read off:

$$v^2 = \frac{1}{\mu_0 \epsilon_0}$$

Taking the positive square root gives the speed of an electromagnetic wave in vacuum:

$$\boldsymbol{c = \frac{1}{\sqrt{\mu_0 \epsilon_0}}}$$

In the unified-wave-speed picture of Chapter 7, every wave speed comes from the square root of an elastic property divided by an inertial property. The same structure repeats here, with the **permeability μ_0** playing the role of an *inertial property* of the vacuum and the **reciprocal of the permittivity $1/\epsilon_0$** playing the role of an *elastic property*. Both are properties of the vacuum itself rather than of a string or a gas. The vacuum is the medium that supports the wave, though the word "medium" must be taken loosely: there is nothing in the vacuum to shake, only fields that propagate through it.
The numerical value of c follows by substituting the modern measured values $\mu_0 = 4\pi \times 10^{-7}$ T·m/A and $\epsilon_0 = 8.854 \times 10^{-12}$ F/m.

$$c = \frac{1}{\sqrt{(4\pi \times 10^{-7})(8.854 \times 10^{-12})}} = 2.998 \times 10^8 \text{ m/s} \approx 3 \times 10^8 \text{ m/s}$$

When Maxwell arrived at this number in 1865 he stared at it. The value he had obtained from purely electromagnetic constants was the same, to within the precision of the measurements then available, as the experimentally known speed of light in vacuum, measured by Hippolyte Fizeau in 1849 and by Léon Foucault in 1862.

The conclusion was unavoidable. *Light is an **electromagnetic wave**.* The visible light that warms a face in the sun, the green of a banana leaf, the blue of the Indian Ocean off Bagamoyo, are oscillating electric and magnetic fields propagating through space at the speed Maxwell's equations predicted. Forty years later Einstein would build his special relativity on the foundation that this speed is the same for every observer; ninety years later quantum electrodynamics would describe the same fields as photons. But Maxwell saw the wave nature of light first, and saw it with nothing more than a pencil, a sheet of paper, and the equations of electromagnetism.

Properties of an electromagnetic wave

A plane electromagnetic wave propagating in vacuum has four essential properties.

1) The wave is transverse

Both the **electric field E** and the **magnetic field B** oscillate in directions perpendicular to the direction of propagation. There is no analogue here of the longitudinal compression that defined sound earlier in this chapter, because there is no medium to compress. The transverseness is intrinsic to the electromagnetic field itself.

2) The two fields are perpendicular to each other

At any point in space and at any instant, E and B are perpendicular both to the direction of propagation and to each other. The three directions, that of E, that of B, and that of propagation, form three mutually perpendicular axes, with the direction of propagation given by the right-hand rule applied from E to B.

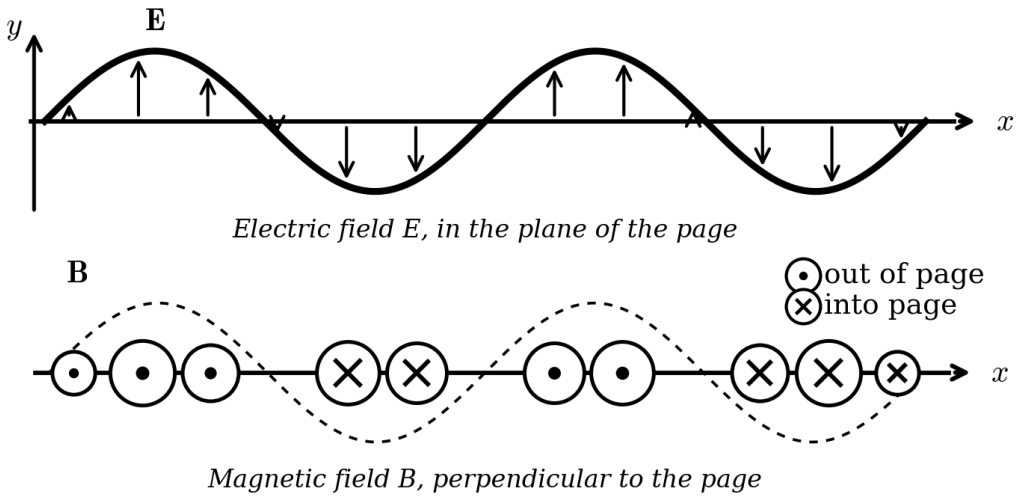


Figure: ***Structure of a plane electromagnetic wave.*** *The electric field E oscillates in the y-direction (in the plane of the page); the magnetic field B oscillates in the z-direction (perpendicular to the page), shown using the standard dot-and-cross symbols. Both fields reach their maxima at the same x and are in phase with each other.*

3) The two fields are in phase

They reach their maxima at the same place and at the same time and pass through zero together, in lock step. This is a consequence of the coupled structure of Maxwell's equations: each field generates a changing component of the other as the wave propagates, and the only consistent solution has the two fields oscillating together.

4) The magnitudes are linked by the wave speed

At every point in space and every instant, the amplitudes of the two fields satisfy

$$E = cB$$

so that a wave with electric-field amplitude $E_0 = 1$ V/m has magnetic-field amplitude $B_0 = \frac{E_0}{c} = 3.33 \times 10^{-9}$ T, a very small number. EM waves in practice are usually described by their electric-field amplitude, because the corresponding magnetic-field amplitude is so small numerically as to be inconvenient to write down.

One immediate consequence of the transverseness of an electromagnetic wave is that the wave can be polarised. The direction in which the electric field oscillates from cycle to cycle is called the polarisation direction of the wave, and a wave whose electric field stays in one plane is said to be plane-polarised. A longitudinal wave, like sound, has no polarisation direction to speak of, because its displacement is along the direction of travel and no transverse choice is to be made. Polarisation is

therefore a property unique to transverse waves, and the full treatment of how it is produced, manipulated, and detected was given in Chapter 8 on physical optics.

Propagation: sine-form description of E and B

The four geometric properties just listed, describe the wave at a single instant of time. They do not yet describe how the wave actually propagates: how the pattern of E and B shifts down the x-axis as time passes, and how a single fixed point in space sees its E and B oscillate up and down with time. The propagation picture is supplied by the sine forms of E and B, which are the explicit time-and-space solutions of the wave equations Maxwell's prediction gave us.

Consider a plane electromagnetic wave travelling in the positive x-direction with angular frequency ω and wavenumber k. The electric field, polarised in the y-direction, is

$$\mathbf{E}(\mathbf{x},\mathbf{t}) = \mathbf{E_0}\,\mathbf{sin}(\mathbf{kx}-\boldsymbol{\omega}\mathbf{t})$$

and the magnetic field, polarised in the z-direction, is

$$\mathbf{B}(\mathbf{x},\mathbf{t}) = \mathbf{B_0}\,\mathbf{sin}(\mathbf{kx}-\boldsymbol{\omega}\mathbf{t})$$

where E_0 and B_0 are the amplitudes of the two fields. Four features of these sine forms deserve attention.

- **Same argument $\boldsymbol{kx-\omega t}$ for both fields.** Both E and B depend on the same combination $kx-\omega t$ of position and time. This is exactly the in-phase property listed earlier: at every x and every t, the argument of the sine is the same for the two fields, so they peak together, pass through zero together, and reach their troughs together. The two oscillations are locked in step, not by accident but by the structure of Maxwell's equations.
- **Two amplitudes, linked by $\boldsymbol{E_0 = cB_0}$.** The amplitudes E_0 and B_0 are not independent; they are tied by the magnitude relation $E = cB$ from the properties list, applied to the peak values. The wave's strength is fully specified by either of E_0 or B_0 alone, and standard practice is to specify E_0 because the corresponding B_0 is so small numerically.
- **The wavenumber $\boldsymbol{k}$ and angular frequency $\boldsymbol{\omega}$.** The wavenumber k, measured in radians per metre, governs how rapidly the field changes with position at a fixed time: a complete cycle takes a distance $\lambda = 2\pi/k$, the wavelength. The angular frequency ω, measured in radians per second, governs how rapidly the field changes with time at a fixed position: a complete cycle takes a time $T = 2\pi/\omega$, the period. The ordinary frequency f and the ordinary wavelength λ are related to ω and k by:

$$\omega = 2\pi f \quad \text{and} \quad k = \frac{2\pi}{\lambda}$$

- **The wave speed is $\boldsymbol{\omega/k}$, and it equals $\boldsymbol{c}$.** Substituting the sine forms into the wave equation derived earlier in this section, $\partial^2 E/\partial t^2 = (1/\mu_0\epsilon_0)\,\partial^2 E/\partial x^2$, the time derivative produces a factor of ω^2 and the spatial derivative a factor of k^2. The wave equation is satisfied provided

$$\omega^2 = \frac{k^2}{\mu_0\epsilon_0}$$

Taking the positive square root and rearranging,

$$\frac{\omega}{k} = \frac{1}{\sqrt{\mu_0\epsilon_0}} = c$$

The ratio ω/k is the speed at which the pattern of the wave moves along the x-axis, and that speed comes out to be exactly the c derived earlier from the coefficient of the wave equation. The same value of c that the wave equation predicted in the abstract is now confirmed concretely by the propagation of the sine pattern itself. The same conclusion is reached from the more familiar pair of relations $c = f\lambda$ and $c = \omega/k$, which the student has met in Chapter 7 for mechanical waves; the new content here is that the value of c comes out to be $1/\sqrt{\mu_0\epsilon_0}$ for any electromagnetic wave, irrespective of its frequency or wavelength.

The student should imagine the sine forms vividly. At a fixed moment in time, E traces out a sinusoid in space along the x-axis (the snapshot the figure shows); as time advances by one period T, the entire sinusoid slides forward by exactly one wavelength λ at speed $c = \lambda/T = \omega/k$. A fixed point in space, like a TV antenna or a retinal cell in a human eye, watches the field at that location oscillate sinusoidally in time with period T; this is the time-domain picture. The two pictures, the spatial snapshot and the time-domain oscillation at one point, are the same wave seen from two different angles.

Speed of light in a material medium

In a transparent material such as glass or water, the speed of an electromagnetic wave is not c. The material's atoms respond to the passing wave by polarising and re-radiating in step with it, and the net effect is to slow the wave down compared with its vacuum speed. The slowdown is captured by a single number, the **refractive index $\boldsymbol{n}$** of the material, defined by:

$$\mathbf{v} = \frac{\mathbf{c}}{\mathbf{n}}$$

The same number n is the one introduced in Chapter 8 in connection with Snell's law of refraction at a boundary between two media. The two definitions are consistent: a wave that slows down on entering a denser material also bends towards the normal at the boundary, and the same ratio n governs both effects.

A few representative values, which a Tanzanian A-level student should recognise on sight, are $n = 1.00$ for vacuum (exactly by definition), $n = 1.00029$ for dry air at sea-level pressure (almost equal to vacuum), $n = 1.33$ for fresh water, $n = 1.50$ for ordinary crown glass, $n = 1.52$ for the dense flint glass used in optical instruments, and $n = 2.42$ for diamond. **Diamond's high refractive index** *is what gives the polished stone its characteristic sparkle, because light entering it slows down dramatically and bounces back and forth inside the stone many times before escaping.*

The refractive index of a material in turn comes from its electromagnetic properties, in exact analogy with $c = 1/\sqrt{\mu_0 \epsilon_0}$ in vacuum. If the material has absolute permeability μ and absolute permittivity ϵ, then

$$v = \frac{1}{\sqrt{\mu \epsilon}},$$

and the refractive index is:

$$n = \frac{c}{v} = \sqrt{\frac{\mu \epsilon}{\mu_0 \epsilon_0}}$$

For most non-magnetic materials μ is very close to μ_0, so the formula simplifies to:

$$n \approx \sqrt{\frac{\epsilon}{\epsilon_0}},$$

and the optical behaviour of the material is determined by its dielectric properties alone.

Energy and intensity of an electromagnetic wave

Just like the mechanical waves of Chapter 7, an EM wave carries energy. The energy is stored partly in the electric field and partly in the magnetic field, both expressed per unit volume. The energy density in the electric field is

$$u_E = \frac{1}{2}\epsilon_0 E^2$$

and the energy density in the magnetic field is

$$u_B = \frac{1}{2\mu_0} B^2$$

In an EM wave the two energy densities are equal at every instant, because $E = cB$ gives $B^2 = E^2/c^2 = \mu_0 \epsilon_0 E^2$, so that $u_B = \left(1/(2\mu_0)\right) \times \mu_0 \epsilon_0 E^2 = 1/2\, \epsilon_0 E^2 = u_E$. The energy is split exactly half-and-half between the electric and magnetic fields. The total energy density is:

$$\mathbf{u = u_E + u_B} = \frac{1}{2}\epsilon_0 E^2 + \frac{1}{2}\epsilon_0 E^2 = \boldsymbol{\epsilon_0 E^2}$$

Averaged over one full cycle of oscillation, the squared field averages to half its peak value, so the time-averaged total energy density is:

$$\bar{u} = \frac{1}{2}\epsilon_0 E_0^2$$

The intensity of the wave is the rate at which energy crosses a unit area perpendicular to the propagation direction. Energy is carried at the wave speed c, so the intensity is

$$I = \bar{u} \cdot c = \frac{1}{2}\epsilon_0 c E_0^2$$

This is the formula a student needs whenever an EM wave's intensity must be related to the amplitude of its electric field.

Foucault's rotating mirror: measuring c on a laboratory bench

Maxwell predicted c from theory; Léon Foucault, in 1862, measured it on a laboratory bench in Paris. His method was clever enough to deserve a place in every undergraduate physics curriculum since, and clean enough that a Form Five student with access to a fast-spinning mirror, a long room, and a graduated screen could replicate it in principle.

The apparatus, illustrated in the following figure, consists of four optical elements. A small light source S emits a thin pencil of light. The light reflects off a small rotating mirror M_R that spins about a vertical axis at a known angular speed ω. The reflected beam travels a long horizontal distance D and reflects off a fixed concave mirror M_F mounted at the far end of the apparatus. The returning beam strikes the rotating mirror a second time. Because the rotating mirror has rotated through a small angle in the time the light took to make its round trip, the second reflection does not retrace the outgoing path. Instead, the returning beam is deflected by a measurable amount onto a screen placed at a known distance r from M_R, and the deflection is read off the screen.

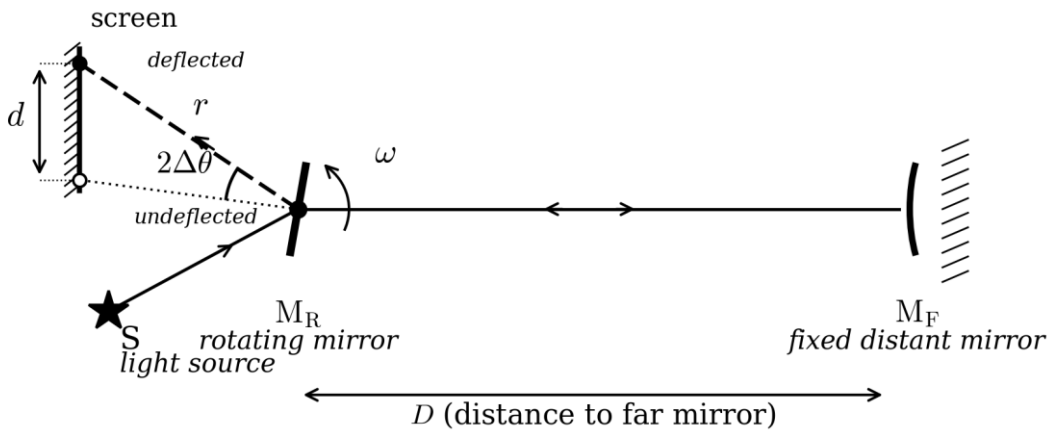


Figure: ***Schematic of Foucault's rotating-mirror apparatus for measuring c.*** *Light from source S reflects off the rotating mirror M_R, travels distance D to the fixed mirror M_F, and returns to M_R. The dotted line shows the path the returning beam would take if M_R had not rotated during the round trip; the dashed line shows the actual path, deflected by an angle $2\Delta\theta$ because M_R has rotated by $\Delta\theta$ during the time the light was in flight. The deflected beam strikes the screen at distance r from M_R, displaced by d from the spot at which the undeflected beam would have arrived.*

The derivation is short. Let the round-trip time for the light be t. Then:

$$t = \frac{2D}{c}$$

During this time the rotating mirror sweeps through an angle,

$$\Delta\theta = \omega t = \frac{2D\omega}{c}$$

By the law of reflection, a rotation of the mirror by $\Delta\theta$ rotates the reflected beam by $2\Delta\theta$. The returning beam is therefore deflected relative to the original outgoing direction by:

$$2\Delta\theta = \frac{4D\omega}{c}$$

The displacement d of the spot on the screen at distance r from M_R, for small angles, is the deflection angle times the screen distance:

$$d = r \times 2\Delta\theta = \frac{4Dr\omega}{c}$$

Solving for c:

$$\mathbf{c = \frac{4Dr\omega}{d}}$$

This is Foucault's formula. The right-hand side contains only quantities the experimenter measures with care on a laboratory bench: the two distances D and r, the angular speed ω of the rotating mirror, and the displacement d of the returning spot on the screen. Foucault's 1862 measurement obtained $c = 2.98 \times 10^8$ m/s, within 0.7 per cent of the modern accepted value, using a rotating mirror driven by a falling weight and a distant fixed mirror 20 metres down the corridor of an observatory.

Foucault's method is one of two classic nineteenth-century laboratory measurements of c. The other is **Fizeau's toothed-wheel method** of 1849, in which a beam of light is chopped into pulses by a rotating cogwheel, sent to a distant mirror, and chopped a second time on its return; the speed of light is recovered from the rotational speed at which the second chop just lets the returning pulse through, given the gap between adjacent teeth and the distance to the distant mirror. We restrict ourselves to Foucault's apparatus here because it is the simpler of the two to derive from first principles, and it is more common.

We now have the six pillars of the electromagnetic chapter. The worked examples that follow exercise four of these pillars in turn, beginning with a direct numerical application of Maxwell's formula and ending with a question that asks the student to compute the electric and magnetic fields of sunlight at the Earth's surface in Tanzania on a clear noon.

BINDER Example 14

(a) Taking the permeability of free space as $\mu_0 = 4\pi \times 10^{-7}\ \mathrm{T \cdot m/A}$ and the permittivity of free space as $\epsilon_0 = 8.854 \times 10^{-12}\ \mathrm{F/m}$, compute the speed of an electromagnetic wave in vacuum to four significant figures.

(b) The refractive index of ordinary crown glass is $n = 1.50$. Compute the speed of light in this glass.

Solution

(a) Apply Maxwell's prediction directly.

$$c = \frac{1}{\sqrt{\mu_0 \epsilon_0}} = \frac{1}{\sqrt{(4\pi \times 10^{-7}\,\mathrm{T\cdot m/A}) \times (8.854 \times 10^{-12}\,\mathrm{F/m})}} = 2.998 \times 10^8\ \mathrm{m/s}$$

To four significant figures the speed of light in vacuum is $c = 2.998 \times 10^8$ m/s.

(b) For glass, $v = c/n$:

$$v = \frac{c}{n} = \frac{2.998 \times 10^8\ \mathrm{m/s}}{1.50} = 1.999 \times 10^8\ \mathrm{m/s}$$

The speed of light in the glass is 1.999×10^8 m/s.

Making sense of the answer: *Two universal constants of electromagnetism, neither of which has anything obvious to do with light, combine in part (a) to give the speed at which light travels. The numerical coincidence is so striking that Maxwell took it as proof that light is itself an electromagnetic phenomenon, twenty-six years before Heinrich Hertz produced electromagnetic waves at radio frequencies and verified that they had every property Maxwell had predicted. The slowdown in glass in part (b) is large enough to matter for any optical instrument: it is the reason a glass prism bends light into a spectrum, the reason a swimming pool looks shallower than it is, and the reason a lens focuses parallel rays to a point.*

Think like a physicist: *Whenever a formula in physics combines two constants from different branches of the subject and reproduces a known result from a third branch, sit up and pay attention. The unification of optics and electromagnetism by Maxwell, of mechanics and gravitation by Newton, of electricity and magnetism by Faraday and Ampère, and of matter and energy by Einstein, are all of this form. Most progress in physics is the unification of two things that were previously thought to be three.*

REAL Example 15

TANESCO uses a long-distance microwave radio link between two relay towers to coordinate power-grid operations. The link operates at a frequency of 6.0GHz.

(a) Compute the wavelength of the microwave signal in air.

(b) Why is microwave preferred over visible light for line-of-sight communication between two relay towers?

(c) Why is microwave not used to illuminate a stage at a concert, even though it would carry the same energy a stage spotlight does?

Solution

(a) Treating the air as effectively a vacuum ($n_{air} \approx 1.000$), the speed of the microwave is essentially $c = 3.0 \times 10^8$ m/s. From the basic wave relation $c = f\lambda$,

$$\lambda = \frac{c}{f} = \frac{3.0 \times 10^8\ \mathrm{m/s}}{6.0 \times 10^9\ \mathrm{Hz}} = 0.050\mathrm{m} = 5.0\mathrm{cm}$$

The signal has a wavelength of 5 centimetres in air, about the width of a smartphone.

(b) Microwave is preferred over visible light for line-of-sight tower-to-tower communication for three connected reasons.

First, the 5cm microwave wavelength is short enough that a dish antenna of practical size (a metre or so across) can focus the signal into a tight, narrow beam aimed precisely at the receiving tower, which keeps the transmitted power on target and reduces interference with neighbouring links. A visible-light beam over the same range would also be directional, but to focus visible light into a narrow beam over kilometres of atmosphere requires a laser, not a passive dish.

Second, microwave at 6GHz penetrates rain, fog, mist, and dust far better than visible light does. The rainy season in Mwanza or Mbeya routinely blocks visible light over a few hundred metres, but a microwave link continues to operate.

Third, the manufacture of high-power microwave transmitters (magnetrons, klystrons, and modern solid-state amplifiers) is mature and economical, whereas high-power visible-light transmitters that could carry equivalent data rates over equivalent distances would be substantially more expensive and less reliable.

(c) Microwave is not used to illuminate a stage because the human eye **does not detect** microwave at all. The retina is sensitive only to electromagnetic waves with wavelengths in the band from about 400nm to 700nm, and 5cm microwave is some hundred thousand times longer than the longest wavelength the eye can register. Replacing a stage spotlight with a microwave transmitter of the same power would leave the stage in pitch darkness, while **heating** the audience uncomfortably. The same energy can be delivered by either wave, but the perception of the energy depends on its wavelength. A successful stage spotlight must lie inside the band the human eye can detect, and that band is the band we call visible light.

Making sense of the answer: *The whole microwave band (1mm to 1m in wavelength) sits between the radio band and the infrared band on the electromagnetic spectrum. At 5cm it sits well inside the part of the band used for telecommunications, broadcast satellite, and weather radar. The Vodacom mobile-phone band at 900MHz uses 33cm wavelengths, longer than the TANESCO link in this question; the Wi-Fi router in a Dar es Salaam student's home at 2.4GHz uses 12.5cm; and a kitchen*

microwave oven at 2.45GHz also uses 12.5cm. The choice of a particular frequency for a particular application is always a compromise between antenna size, atmospheric transparency, and bandwidth, and the engineer's job is to find the band that wins on all three.

Think like a physicist: *Every part of the electromagnetic spectrum can carry energy, can be modulated to carry information, and can be focused by a properly designed antenna. What changes from band to band is which technological problems become easy and which become hard. A radio engineer designing a link of any kind asks four questions before choosing a frequency: how directional must the beam be (which sets the smallest workable antenna at any given size), how much rain and fog must it punch through (which sets the highest workable frequency), what data rate is required (which sets the lowest workable frequency through the bandwidth-per-cycle argument), and how much licensing the regulator will charge. The Tanzania Communications Regulatory Authority (TCRA) controls the answer to the fourth.*

REAL Example 16

A physics teacher at Mzumbe Secondary School sets up a Foucault apparatus on the long corridor outside the laboratory. A small rotating mirror, spinning at 400 revolutions per second about a vertical axis, sends a thin beam of light to a fixed concave mirror at a distance of 35m down the corridor. The light returns to the rotating mirror and is reflected onto a screen placed at a perpendicular distance of 5.0m from the rotating mirror. The displacement of the returning spot on the screen, relative to its position when the rotating mirror is stationary, is measured to be 5.87mm.

(a) Use this data to compute the speed of light in air, and compare your result with the modern accepted value.

(b) Without redoing the experiment, predict the displacement d that would be observed if the rotation rate were doubled to 800 revolutions per second, with everything else unchanged. Justify your prediction with reference to the formula derived in this section.

Solution

(a) Foucault's formula is $c = 4Dr\omega/d$, where ω is the angular speed of the rotating mirror in radians per second. The rotation rate is given as 400 revolutions per second, and each revolution sweeps an angle of 2π radians, so:

$$\omega = 2\pi f = 2\pi \times 400\ \text{Hz} = 2513\ \text{rad/s}$$

Substitute the four measured quantities into Foucault's formula:

$$c = \frac{4Dr\omega}{d} = \frac{4 \times 35\ \text{m} \times 5.0\ \text{m} \times 2513\ \text{rad/s}}{5.87 \times 10^{-3}\ \text{m}} = 2.997 \times 10^{8}\ \text{m/s}$$

The measured speed of light is $c = 2.997 \times 10^{8}$ m/s, which agrees with the modern accepted value of 2.998×10^{8} m/s to four significant figures. The small discrepancy of about 0.03 per cent could come from rounding the rotation rate to 400Hz exactly, from a small misalignment of the screen, or from the slightly slower speed of light in air than in vacuum (in air, c is about 0.03 per cent lower than in vacuum, which would account for the entire discrepancy on its own).

(b) The Foucault formula in its forward form is $d = 4Dr\omega/c$. Holding D, r, and c fixed, the displacement d is directly proportional to ω:

$$d \propto \omega$$

If ω doubles, d doubles. The predicted displacement at 800 revolutions per second is therefore

$$d_{\text{new}} = 2 \times 5.87\text{mm} = 11.7\text{mm}$$

The experimenter would observe the spot move twice as far on the screen as in part (a).

Making sense of the answer: *That a corridor of 35m, a mirror spinning at 400 revolutions per second, and a deflection measured in millimetres, can deliver the speed of light to four significant figures, is one of the great miracles of nineteenth-century experimental physics. The geometric amplification is what makes it work: a tiny rotation angle of the mirror ($2\Delta\theta$ in this experiment is about 1.2×10^{-3} radians, or about a fifteenth of a degree) becomes a millimetres-large displacement on the screen because the screen is several metres away from the mirror. Tiny angles become measurable distances when the lever arm is long enough. Part (b) of the question turns the formula around and uses it predictively: once the relationship between input and output is established, the experimenter can predict outputs for new inputs without measuring afresh.*

Think like a physicist: *Foucault's apparatus is an early example of what experimental physicists call a leverage arm: an arrangement in which a small quantity of interest is amplified into a much larger quantity that can be measured directly. The same principle appears again and again throughout experimental physics: the challenge is always to arrange the system so that an extremely small effect becomes large enough to detect before it disappears into background noise.*

HOT Example 17

The intensity of sunlight at the Earth's surface in Tanzania on a clear day at noon is approximately $1000\ \text{W/m}^2$. Treating the sunlight as a plane electromagnetic wave with a single dominant frequency, find:

(a) the amplitude E_0 of the electric field at the surface;

(b) the amplitude B_0 of the corresponding magnetic field;

(c) the total electromagnetic energy striking a one-square-metre solar panel mounted perpendicular to the sun's rays in one hour of clear noon.

Take $\epsilon_0 = 8.854 \times 10^{-12}$ F/m and $c = 3.00 \times 10^8$ m/s.

Solution

(a) Apply the intensity formula $I = 1/2\,\epsilon_0 c E_0^2$. Solving for E_0,

$$E_0^2 = \frac{2I}{\epsilon_0 c}$$

$$E_0 = \sqrt{\frac{2I}{\epsilon_0 c}} = \sqrt{\frac{2 \times 1000\ \text{W/m}^2}{(8.854 \times 10^{-12}\ \text{F/m}) \times (3.00 \times 10^8\ \text{m/s})}} = 868\ \text{V/m}$$

The electric-field amplitude of sunlight at the surface is 868V/m.

(b) The corresponding magnetic-field amplitude follows from $B_0 = E_0/c$:

$$B_0 = \frac{E_0}{c} = \frac{868\ \text{V/m}}{3.00 \times 10^8\ \text{m/s}} = 2.89 \times 10^{-6}\ \text{T} \approx 2.9\mu\text{T}$$

The magnetic-field amplitude is therefore about 2.9 microtesla, comparable to the Earth's magnetic field but oscillating at the frequency of visible light.

(c) The total energy striking a one-square-metre solar panel in one hour is the intensity times the area times the time:

$$\text{Energy} = I \times A \times t = \left(1000\text{W/m}^2\right) \times (1\text{m}^2) \times (3600\text{s}) = 3.60 \times 10^6\ \text{J} = 3.6\ \text{MJ}$$

A one-square-metre solar panel exposed to clear noon sunlight in Dar es Salaam receives 3.6 megajoules of energy in one hour. At a typical solar-panel conversion efficiency of 20 per cent, this corresponds to about 720 kilojoules of useful electrical energy, or roughly 0.2 kilowatt-hours per square metre per hour of clear noon. A 5-square-metre roof installation, exposed to the equivalent of 5 hours of clear noon per day, would generate about 5 kilowatt-hours per day, which is roughly the daily electricity needs of a small Tanzanian household.

Making sense of the answer: *The electric-field amplitude of sunlight is about 870V/m, which sounds large until one remembers that the dielectric strength of dry air is about 3 million V/m, so sunlight is operating at less than one part in 3000 of the field strength needed to ionise air. The magnetic-field amplitude of 2.9 microtesla is comparable to Earth's magnetic field, but the sunlight oscillates at frequencies near* 5×10^{14} *Hz, while Earth's field is static. The 3.6 megajoules of energy in one hour is enough to heat about 10 kilograms of water by 86 degrees, which is why a black-painted water tank on a tin roof in Dodoma produces hot water all day long without any external power.*

Think like a physicist: *Every wave that carries energy can be characterised by the amplitude of the quantity that oscillates: pressure for sound, displacement for a wave on a string, electric field for light. The intensity formula always has the same structure: half times something times something times amplitude squared, where the two "somethings" are the medium's response constants. For light it is* $1/2\,\epsilon_0 c E_0^2$. *For sound in a fluid it is* $1/2\,\rho v \omega^2 s_0^2$ *where* s_0 *is the displacement amplitude. The pattern repeats across the entire wave physics of the universe.*

We have now answered, in physical terms, every question the chapter set for itself about electromagnetic waves. But the family of electromagnetic waves is broad, and the visible light that has occupied our discussion so far is only one narrow band in a spectrum that extends from radio waves at one extreme to gamma rays at the other. The next section catalogues the electromagnetic family by frequency, describes how each band is produced and detected, and lists the applications by which each band touches the daily life of an A-level student, from the morning radio bulletin to the evening television and from the X-ray of a broken arm to the microwave that warms last night's leftovers.

THE ELECTROMAGNETIC SPECTRUM

The previous section closed with the observation that visible light is a single narrow band inside a much broader electromagnetic family. The family runs from radio waves of kilometre wavelength at one end to gamma rays of sub-picometre wavelength at the other end, and every member shares the same fundamental properties: a transverse oscillation of electric and magnetic fields, propagation through vacuum at speed c, the in-phase relationship $E = cB$ between amplitudes, and the intensity formula $I = 1/2\,\epsilon_0 c E_0^2$. What changes from one member to the next is wavelength, and with wavelength every practical consequence: how the wave is produced, how it is detected, what it does on contact with matter, and what technology, medicine, and engineering have built around it. This section catalogues the seven bands the syllabus recognises, gives a Tanzanian instance for each, and ends with one worked example on the most familiar microwave instrument the student has met, the kitchen microwave oven at 2.45GHz.

The spectrum at a glance

The following figure shows the entire electromagnetic spectrum on a single logarithmic frequency axis. Frequency increases from left to right, by two powers of ten per tick, from 10^4 Hz on the far left (very-low-frequency radio) to 10^{22} Hz on the far right (high-energy gamma rays from nuclear processes). Wavelength, plotted along the top of the same axis, decreases correspondingly from 30km on the far left to 30fm (femtometres) on the far right, an overall span of eighteen orders of

magnitude. Three Tanzanian-context frequencies are marked above the bar: the TBC FM band at 100MHz, the kitchen microwave oven at 2.45GHz, and a medical X-ray at the Muhimbili National Hospital at about 0.1nm. The visible band is so narrow on this scale that it is shown by hatching and a leader line, with a magnified inset below the main bar showing the seven colours of the visible rainbow.

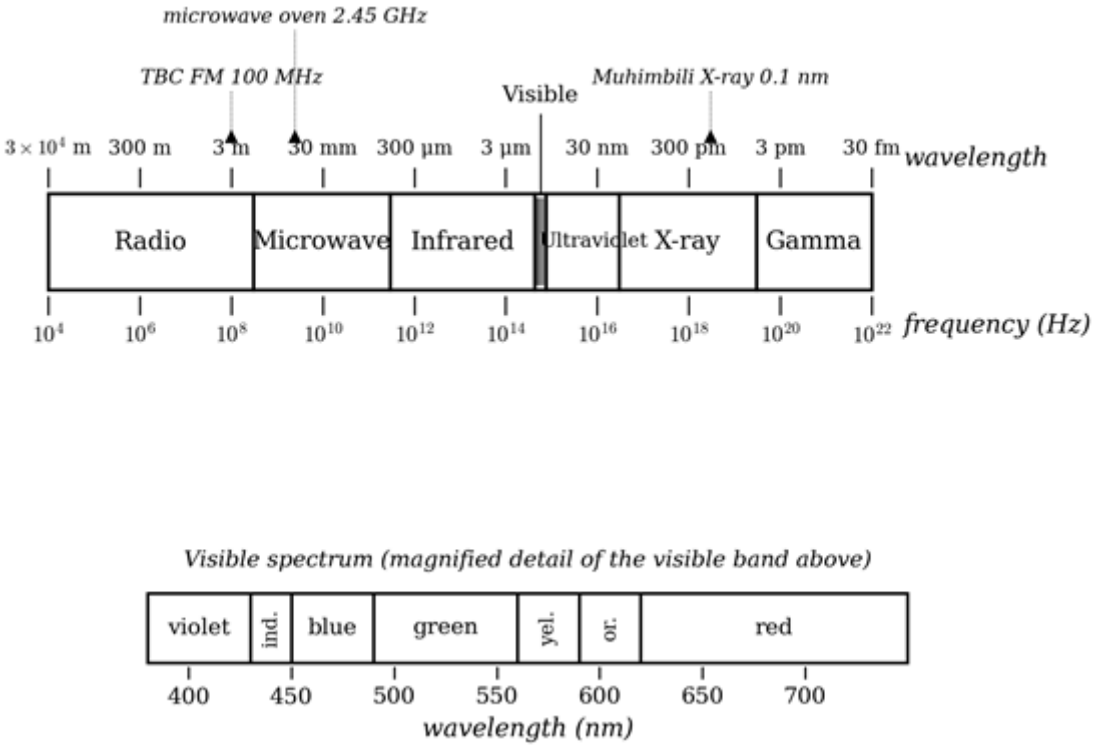


Figure: ***The electromagnetic spectrum.*** *Top: the full spectrum from radio to gamma, on a logarithmic axis spanning eighteen orders of magnitude in frequency. Three Tanzanian-context frequencies are marked above the bar. Bottom: a magnified inset of the visible band (380–750nm), showing the seven colours violet, indigo, blue, green, yellow, orange, red.*

The student should pause and stare at this figure for a moment. The visible light that drives every act of human sight, from reading this page to watching a sunset over the Indian Ocean, occupies less than one twentieth of one decade on a logarithmic scale eighteen decades wide. Everything else, from the FM radio in a car's dashboard to the cobalt-60 source at the Ocean Road Cancer Institute, is invisible to the eye yet identical in nature to the light the eye does see. Maxwell's prediction unified them all, and the seven bands below are a catalogue of the same wave, partitioned only by what humans have learnt to do with it at each scale.

Radio waves occupy the longest wavelengths, from kilometres at the very-low-frequency end down to about one metre at the boundary with microwave. They are produced by accelerating electrons along the length of a metal antenna, driven by an LC oscillator circuit at the desired frequency. They are detected by a receiving antenna, tuned by a resonant LC circuit at the listener's end, and demodulated by a detector. In Tanzania, AM radio operates in the kilohertz band, FM radio (including the Tanzania Broadcasting Corporation, TBC, and Clouds FM) operates around 100MHz, and short-wave radio reaches the rest of East Africa from transmitters in Dar es Salaam. The same band carries television broadcast, mobile-phone signals at the low end of the microwave boundary, and amateur radio.

Microwaves occupy wavelengths from 1m down to 1mm, equivalently frequencies from 300MHz to 300GHz. They are produced by specialised oscillators such as klystrons and magnetrons (the latter is the source inside every kitchen microwave oven) or by modern solid-state circuits. They are detected by crystal diodes paired with horn antennas or parabolic dishes. Tanzanian instances are abundant: Vodacom and Airtel mobile-phone networks operate at 900, 1800, 2100, and 2600MHz; the SEACOM undersea optical cable lands microwave-relayed traffic in Dar es Salaam; the TANESCO long-distance microwave radio link of REAL Example 15 in the previous section operates at 6GHz; satellite television uses the 4GHz C-band; and the kitchen microwave oven, the subject of the worked example at the end of this section, operates at 2.45GHz.

Infrared occupies wavelengths from 1mm down to about 700nm at the boundary with visible red, equivalently frequencies from 300GHz up to about 430THz. Infrared is produced by every warm object in the universe simply by virtue of being warm; the spectrum a body emits is determined by its temperature, and the warmer the body the shorter the peak wavelength. A human body at 37°C radiates predominantly in the infrared band centred around 10µm. Infrared is detected by bolometers (which measure the heat the radiation deposits), thermopiles, and the silicon-array infrared cameras now built into every smartphone selfie camera in altered form. Tanzanian instances: every television remote control uses an 850nm infrared LED to send a command to the television; airport thermal scanners use infrared cameras to screen arriving passengers for fever;

night-vision goggles, used by anti-poaching units patrolling the Selous Game Reserve, see the warm bodies of poachers and wildlife against the cool background of bush at night.

Visible light occupies the narrow band from 700nm (red) down to 400nm (violet), equivalently a frequency window from 430THz to 750THz. The seven colours of the rainbow violet, indigo, blue, green, yellow, orange, red sit inside this window, each occupying a slice perhaps 30 to 50nm wide. Visible light is produced by electronic transitions in atoms, including the heated filament of an incandescent bulb (giving a continuous spectrum of blackbody emission), the discrete spectral lines of a sodium-vapour streetlamp, the white light of a light-emitting diode, and the broadband emission of the Sun. The eye is the most sensitive detector of visible light known, with the rod cells of the dark-adapted retina capable of registering individual photons; photographic film, semiconductor photodiodes, and CCD arrays in digital cameras are the engineered equivalents. Tanzanian instances are everywhere: sight itself, photography, fibre-optic telecommunications at the special wavelengths 1310nm and 1550nm (technically in the near infrared, but the same physics: the SEACOM cable that lands at Dar es Salaam carries internet traffic at these wavelengths). The seven colours of the rainbow that follow a Dar es Salaam thunderstorm are this same visible band, dispersed by the prism action of the raindrops.

Ultraviolet occupies the band from 400nm down to about 10nm, equivalently 750THz up to 3×10^{16} Hz. Ultraviolet is produced by hotter sources than visible light: the mercury arc lamp, the very hot surface of the Sun, and electronic transitions involving larger energy gaps. It is detected by fluorescent screens (which absorb ultraviolet and re-emit it as visible light, the same physics that makes a fluorescent banknote feature glow under a security checker), and by specialised photographic film. The Sun radiates substantial ultraviolet, most of which is absorbed by the ozone layer of the upper atmosphere; the small fraction that reaches the surface of Tanzania is responsible both for the tan a holidaymaker acquires on the beaches of Zanzibar and for the vitamin D the same holidaymaker's skin synthesises in response. The harder ultraviolet ionises atoms and can damage DNA, which is why excessive sun exposure carries a cancer risk and why ultraviolet lamps are used to sterilise surgical instruments and drinking water.

X-rays occupy the band from 10nm down to 0.01nm, equivalently 3×10^{16} Hz up to 3×10^{19} Hz. X-rays are produced by inner-shell electron transitions in heavy atoms and by bremsstrahlung (literally "braking radiation") emitted when a fast electron is suddenly decelerated by hitting a metal target; this is the mechanism inside every medical X-ray tube. They are detected by photographic film, by image intensifiers in modern hospital systems, and by ionisation chambers. The wavelength of an X-ray is comparable to the spacing of atoms in a crystal, which is why X-ray diffraction is the standard tool for finding atomic positions in solids. Tanzanian instances: every chest X-ray taken at Muhimbili National Hospital, every dental X-ray at a clinic in Mwanza, every airport luggage scanner at Julius Nyerere International Airport that checks for contraband under bags.

Gamma rays occupy the shortest wavelengths, below 0.01nm and equivalently above 3×10^{19} Hz. They are produced by nuclear transitions, by the annihilation of an electron with a positron, and by cosmic-ray showers in the upper atmosphere. They are detected by Geiger counters, scintillation detectors, and modern semiconductor gamma cameras. Gamma rays carry enough energy per photon to penetrate substantial thicknesses of lead, which is why gamma sources are stored in heavy lead-lined containers and why gamma therapy is among the most carefully controlled procedures in medicine. The Ocean Road Cancer Institute in Dar es Salaam uses a cobalt-60 gamma source to deliver radiation therapy to cancer patients; the same isotope is used worldwide to sterilise surgical equipment and food. The natural background of gamma rays in Tanzania, from radon in soil and cosmic rays from space, is small but never zero, and adds a few millisieverts per year to every Tanzanian's lifetime radiation exposure.

Each band above is the same wave, propagating at the same speed c, differing only in wavelength. The choice of which technological apparatus to use, which biological effect to expect, and which name to assign, depends on the wavelength alone. Once the wavelength is given, the band is fixed, the production and detection methods follow, and the applications fall into place. The worked example below applies this picture to the most familiar microwave appliance in any Tanzanian kitchen, and asks three questions every student should be able to answer.

REAL Example 18

A kitchen microwave oven operates at a frequency of 2.45GHz, the standard frequency for domestic microwave ovens worldwide.

(a) Compute the wavelength of the microwave radiation inside the oven cavity, assuming the cavity is essentially empty (filled with air, with the small water-containing meal occupying a small fraction of the cavity volume).

(b) Explain why 2.45GHz, rather than some other microwave frequency, is used in domestic microwave ovens.

(c) Inside the oven cavity, standing electromagnetic waves form between the conducting metal walls. Estimate the spacing between adjacent nodes of the standing wave, and explain why the oven includes a rotating turntable.

Solution

(a) Using $c = f\lambda$ with $c = 3.0 \times 10^8$ m/s and $f = 2.45 \times 10^9$ Hz,

$$\lambda = \frac{c}{f} = \frac{3.0 \times 10^8 \text{ m/s}}{2.45 \times 10^9 \text{ Hz}} = 0.122 \text{ m} = 12.2 \text{ cm}$$

The microwave inside the oven has a wavelength of about 12cm, comparable to the width of a slice of bread.

(b) Two reasons drive the choice of 2.45GHz, neither involving the resonant frequency of any water molecule.

First, 2.45GHz lies inside an internationally protected slice of the radio spectrum called the **ISM band** (Industrial, Scientific, Medical), which is reserved by international agreement for free, unlicensed use by domestic and industrial appliances that radiate electromagnetic energy. Designing microwave ovens for the ISM band means manufacturers can sell ovens worldwide without negotiating a separate broadcast licence in every country.

Second, water absorbs microwave radiation over a broad range of frequencies, with absorption increasing as the frequency rises. At very high microwave frequencies, near 20GHz, water absorbs so strongly that the radiation penetrates only the outer few millimetres of the food and the inside of the food stays cold; at very low microwave frequencies, near 100MHz, water absorbs so weakly that the radiation passes through the food without depositing significant energy. The frequency 2.45GHz is a compromise: water absorbs enough that a substantial fraction of the radiated power is deposited inside the food, but not so much that the deposition is confined to a thin surface layer. The penetration depth at 2.45GHz is several centimetres, well matched to the typical size of a meal.

It is worth correcting a textbook myth at this point. Many physics books claim that 2.45GHz is the "resonant frequency" of the water molecule, suggesting some sharp peak in absorption at exactly this frequency. This is **not** true. Liquid water has no sharp resonance in the microwave band; instead, it has a broad rotational and dipolar relaxation response that peaks near 20GHz and falls smoothly to zero at much lower frequencies. 2.45GHz sits well below the peak, in the low-frequency tail of that broad band. The choice of 2.45GHz is an engineering compromise between ISM-band availability and useful penetration depth, not a tuning to any resonance.

(c) Inside the oven cavity, the microwave reflects off the metal walls, and the incident and reflected waves combine to form a standing wave with nodes (zero electric field) at fixed positions and antinodes (maximum electric field) between them. As established in the section on resonance in pipes, the spacing between adjacent nodes of a standing wave is half the wavelength of the travelling wave:

$$\text{node spacing} = \frac{\lambda}{2} = \frac{0.122 \text{ m}}{2} = 0.061 \text{ m} \approx 6 \text{ cm}$$

The cavity therefore has hot spots (at the antinodes, where the electric field is large and the heating intense) separated by cold spots (at the nodes, where the field is zero and no heating occurs) approximately 6cm apart. A piece of food sitting in a single position in the oven would have parts inside hot spots and parts inside cold spots, and would heat unevenly. The standard engineering fix is to install a rotating turntable, which spins the food slowly during cooking so that every part of the food spends roughly equal time in hot and cold spots, and the heating averages out across the meal. A microwave oven without a turntable would cook food in a polka-dot pattern, a fact that can be demonstrated with chocolate or marshmallows in any oven by switching the turntable off and observing where the melt-spots appear.

Making sense of the answer: *Every part of this question is grounded in physics already established in this chapter. Part (a) is the basic wave relation* $c = f\lambda$*, the same relation that gave the wavelength of the TANESCO 6GHz microwave in REAL Example 15. Part (b) corrects a common physics misconception by appealing to penetration depth. Part (c) reuses the standing-wave node spacing rule* $\lambda/2$ *from earlier in this chapter on resonance in pipes, applied to a closed metal cavity instead of an open pipe. A microwave oven turns out to be one of the most thoroughly Form-Five-physics-explainable appliances in a Tanzanian kitchen.*

Think like a physicist: *When a physics question describes a real device and asks why it is built the way it is, the answer is almost never "because that is the resonant frequency." Engineers think in terms of compromises across several constraints: cost, availability of components, regulatory permission, performance, safety, and ease of manufacture. The 2.45GHz choice for microwave ovens is a textbook example of an engineering compromise being mis-told as a physics resonance by simplification. Whenever you read a one-line explanation of a real engineered device, ask what compromise it actually represents. The honest physics answer is usually more interesting than the textbook simplification.*

This section has catalogued the entire electromagnetic family by frequency, named each band, identified the typical production and detection methods, and grounded each band in at least one Tanzanian context. The next section turns from the abstract spectrum back to a concrete acoustic question that Chapter 9 has been carrying since its opening dialogue: *when an ambulance passes the listener at speed, sounding its siren, what exactly does the siren's frequency do, and why?* The answer is the Doppler effect, and the next section will derive it from first principles for sound, before applying the same principle to electromagnetic waves and to the redshift of distant galaxies in the section after that.

THE DOPPLER EFFECT FOR SOUND

The opening dialogue of this chapter ended with Mr. Akilikubwa writing a single equation on the blackboard at Jangwani Secondary School,

$$f' = f \times \frac{v \pm u_o}{v \mp u_s}$$

and promising that the next seventy pages would unpack it. The previous four sections have built the toolkit. Sound is a longitudinal mechanical wave in air at a speed of about 340m/s; the speed is set by the medium alone, not by the source; air columns ring at their natural frequencies; the ear hears intensity logarithmically and the eye sees a band of electromagnetic radiation that obeys, in vacuum, exactly the same wave equation. With the toolkit in hand, the chapter must now keep its promise and explain Mr. Akilikubwa's equation, case by case, from the principle that decides it. It must also answer the question Kipute left hanging on the table at the end of the introduction: when the ambulance siren produces the same number of beeps per second whether the vehicle is parked outside the school gate or racing to Muhimbili, *what exactly is the universe doing to make the note rise as it comes and drop as it goes?*

The answer is the **Doppler effect**, named after the Austrian physicist Christian Doppler who described it for sound and for starlight in a paper of 1842. The principle behind every form of the effect, from the ambulance siren on United Nations Road to the redshift of a galaxy in the constellation of Centaurus, can be stated in one statement. Read it once, slowly, before going on.

*In the rest frame of the medium, the wave speed is fixed by the medium alone. Source motion changes the **wavelength** the source lays down behind it. Observer motion changes the **encounter rate** at which the observer meets the wavefronts. The apparent frequency is the wave speed seen by the observer divided by the wavelength as set by the source.*

That statement is the entire physics of the Doppler effect. Every formula in this section follows from it by careful bookkeeping. We will derive the five primary cases, list the three no-Doppler cases (where source motion and observer motion conspire to leave the apparent frequency unchanged), then return to Mr. Akilikubwa's equation and show how the same single formula contains all of them.

The source moves, the observer is at rest

Imagine a stationary source S of sound in still air, emitting a steady tone of frequency f. The source plants a new wavefront into the air every $T = 1/f$ seconds. Each wavefront then propagates outwards through the air at speed v, expanding as a sphere centred on the point where the source emitted it. Adjacent wavefronts are separated by one wavelength,

$$\lambda = v\,T = \frac{v}{f}$$

so that exactly f wavefronts pass any stationary point in the air every second. The picture is symmetric: the wavefronts form a set of concentric circles around the source, in any plane through the source.

Now let the source move. Let the source travel in a straight line through the air at constant speed u_s, while still emitting wavefronts at the same rate f per second. The source is no longer at the same place when it emits the second wavefront as it was when it emitted the first. In the time interval T between two successive emissions, the source has moved a distance $u_s T$ in the direction of its motion. The second wavefront is therefore emitted from a point that is closer to any observer directly in front of the source by exactly that distance.

The geometric consequence is immediate. In front of the source, the wavefronts are closer together than they would have been from a stationary source. Behind the source, they are farther apart. The wave speed in the air is unchanged at v (the medium decides the speed; the source's motion does not), but the wavelength in front is

$$\lambda' = \lambda - u_s T = \frac{v}{f} - \frac{u_s}{f} = \frac{v - u_s}{f}$$

and the wavelength behind is

$$\lambda' = \lambda + u_s T = \frac{v + u_s}{f}$$

A stationary observer directly ahead of the approaching source therefore meets wavefronts of wavelength $(v - u_s)/f$, propagating at the air speed v, and so hears the apparent frequency:

$$\mathbf{f' = \frac{v}{\lambda'} = \frac{v}{(v - u_s)/f} = f \times \frac{v}{v - u_s}}$$

A stationary observer directly behind the receding source meets wavefronts of wavelength $(v + u_s)/f$, propagating at the same air speed, and hears:

$$\mathbf{f' = f \times \frac{v}{v + u_s}}$$

In words: *a source moving towards an observer raises the apparent frequency (the denominator shrinks), and a source moving away from an observer lowers it (the denominator grows).* The siren of the ambulance racing from Kariakoo to Muhimbili rises in pitch as it approaches the Jangwani classroom and drops as it recedes towards the hospital, by exactly this mechanism. The "bunching of wavefronts in front" that Khadija reached for in the chapter's opening dialogue was the correct intuition. What Kipute pressed her for, and what Khadija could not yet quite put into words, was the bookkeeping: *the source moves between two successive emissions, and the wavefront the source has just emitted is therefore born from a point closer to the observer than the previous one was.*

The observer moves, the source is at rest

Now reverse the picture. Let the source be at rest in still air, emitting wavefronts at rate f per second; the wavefronts have wavelength $\lambda = v/f$ and propagate outwards at the air speed v in all directions. The observer, no longer stationary, walks towards the source at speed u_o.

The wavelength of the wave in the air is what the source laid down, $\lambda = v/f$, and the source's stationary condition means that wavelength is the same in every direction. What changes is the speed at which the observer meets successive wavefronts. The observer moving towards the source at u_o sees each wavefront approaching her at $v + u_o$ (the wavefront's speed in the air plus the observer's speed through the air, both measured in the same direction). The rate at which she crosses wavefronts is the encounter speed divided by the spacing,

$$\mathbf{f' = \frac{v + u_o}{\lambda} = \frac{v + u_o}{v/f} = f \times \frac{v + u_o}{v}}$$

If the observer instead walks away from the source at speed u_o, each wavefront catches her at the slower speed $v - u_o$, and she meets them at the lower rate:

$$\mathbf{f' = f \times \frac{v - u_o}{v}}$$

An observer moving towards a stationary source therefore hears a raised frequency, and an observer moving away from a stationary source hears a lowered frequency, by the same physical sign as for a moving source. But the algebra is now different: the moving observer changes the **encounter rate** rather than the wavelength.

The general formula: both source and observer in motion

When both source and observer move through the air, both effects operate simultaneously. The source, moving at u_s relative to the air, alters the wavelength it lays down; the observer, moving at u_o relative to the air, alters the encounter speed at which she meets those altered wavefronts. To arrive at the general formula, we first re-state each of the two component effects on its own, and then combine them step by step.

Source motion alters the wavelength. From the source-moving subsection, the wavelength the source plants in the air, measured in the rest frame of the air, is:

$$\lambda' = \frac{v - u_s}{f} \quad \text{when the source moves towards the observer}$$

$$\lambda' = \frac{v + u_s}{f} \quad \text{when the source moves away from the observer}$$

These two cases can be written as a single expression using a paired-sign symbol,

$$\lambda' = \frac{v \mp u_s}{f}$$

in which the **upper sign (the minus)** applies when the source moves *towards* the observer, shortening the wavelength, and the **lower sign (the plus)** applies when the source moves *away* from the observer, lengthening it.

Observer motion alters the encounter speed. From the observer-moving subsection, the speed at which the observer meets successive wavefronts in the air is:

$$v' = v + u_o \quad \text{when the observer moves towards the source}$$

$$v' = v - u_o \quad \text{when the observer moves away from the source}$$

Combining the two cases into a single expression,

$$v' = v \pm u_o$$

in which the **upper sign (the plus)** applies when the observer moves *towards* the source, raising the encounter speed, and the **lower sign (the minus)** applies when the observer moves *away* from the source, lowering it.

Combining the two effects: The apparent frequency the observer hears is the encounter speed divided by the wavelength, both measured in the rest frame of the air:

$$\mathbf{f' = \frac{v'}{\lambda'} = \frac{v \pm u_o}{(v \mp u_s)/f} = f \times \frac{v \pm u_o}{v \mp u_s}}$$

This is the equation Mr. Akilikubwa wrote on the Jangwani blackboard, derived now from the principle that opened this section.

How to apply the sign convention: The notation $\pm$ and $\mp$ encodes two independent sign choices, one for the observer in the numerator and one for the source in the denominator. The two choices are read separately, from the motion of each.

In the **numerator**, take the **plus** sign if the observer moves *towards* the source, and the **minus** sign if the observer moves *away* from the source. In the **denominator**, take the **minus** sign if the source moves *towards* the observer, and the **plus** sign if the source moves *away* from the observer. The single mnemonic that captures both rules is: *towards* always pushes the apparent frequency upwards (whether it is the source or the observer doing the moving), and *away* always pushes it down.

The choice of $\pm$ in the numerator and $\mp$ in the denominator is the formula's quiet way of reminding the student that the two halves of the bar are inverted with respect to each other. When the numerator is positive ($v + u_o$, observer approaching), the

denominator is negative ($v - u_s$, source approaching); the two choices then push the apparent frequency upwards together. When the numerator is negative ($v - u_o$, observer receding), the denominator is positive ($v + u_s$, source receding); the two choices then push the frequency downwards together. The four physical scenarios are shown in the table below.

Observer motion	**Source motion**	**Numerator**	**Denominator**	**Effect on f'**
towards	towards	$v + u_o$	$v - u_s$	highest possible
towards	away	$v + u_o$	$v + u_s$	depends on speeds
away	towards	$v - u_o$	$v - u_s$	depends on speeds
away	away	$v - u_o$	$v + u_s$	lowest possible

For the two "same direction" cases (both towards each other, or both away from each other) the numerator and denominator pull the apparent frequency in the same direction and the answer is decided in sign. For the two "mixed" cases (one towards, one away), the two effects compete and the answer depends on the relative size of u_s and u_o.

Five further features of the general formula deserve emphasis.

1) The source and observer speeds are measured against the medium, not against each other

The Doppler effect for sound is not a function of relative velocity alone; it is a function of source speed through the air and observer speed through the air, separately. Two cases that look the same from a "relative motion" perspective give different apparent frequencies. (This will become important in REAL Example 21, part (c).) The medium decides everything in the rest frame of the medium.

2) The wave speed v that appears in the formula is the wave speed in the rest frame of the medium

This is the same speed that featured in $c = f\lambda$ throughout Chapter 7. In still air at 20°C that speed is approximately 340m/s. If the air itself moves (a wind blows), the wave speed in the ground frame changes, and a correction is needed; we will examine the wind correction later.

3) The formula applies along the source-to-observer line

If the source moves at an angle to that line, the relevant velocity is the **component along the line**, $u_s\cos\theta$. The same applies to the observer. We treat this oblique case in worked Example 22.

4) The formula is exact only for source and observer speeds well below the wave speed, $u_s, u_o \ll v$

At source speeds approaching v (the speed of sound), the wavefronts pile up into a shock wave and the formula breaks down; at $u_s = v$ the formula's denominator vanishes and the apparent frequency would diverge. This regime, the "sonic boom" of a supersonic aircraft, is qualitatively different physics and is not part of this section.

5) The formula reduces to its four simpler limiting cases by setting one of u_s, u_o to zero, and recovers $f' = f$ in the trivial case $u_s = u_o = 0$

A useful sanity check on any calculation is to set the wrong speed to zero and confirm that the remaining formula matches the four sub-cases of the previous two subsections.

The three no-Doppler cases

You may sometimes encounter a situation involving motion and wonder whether a Doppler shift will be heard. The three principal scenarios in which there is **no** Doppler shift, in spite of motion being present somewhere, are these.

1) Both source and observer at rest in the medium

The trivial case. Wavefronts are emitted at the source's rate, propagate at the medium's speed, and meet the observer at exactly the same rate. $u_s = u_o = 0$ and the formula gives $f' = f$.

2) Source and observer moving at the same velocity through the medium

The geometric consequence is that in the rest frame of either of them, both are stationary; the medium is what moves, and the formulas of the previous subsections, derived in the medium's rest frame, give a non-trivial apparent frequency only when source and observer move differently through the medium. The cleanest example is a passenger sitting inside a moving daladala that is sounding its horn: the source (the horn) and the observer (the passenger) both move with the daladala at the same speed through the air, so neither experiences a Doppler shift relative to the other. The passenger inside the daladala hears the horn at exactly its source frequency, regardless of how fast the daladala travels. We exercise this insight in REAL Example 20, part (c).

3) Source moving along a circle of which the observer is the centre

The source's distance from the observer is constant; the wavefronts the source emits travel outwards along radii of equal length and arrive at the observer at the same rate they are emitted. There is no Doppler shift. (Equivalently, the source's velocity at every instant is perpendicular to the source-to-observer line, so the component along that line is zero, and the oblique formula gives $f' = f$.)

A common student trap is the case of an observer running across a railway platform parallel to a train that is whistling and standing still. The observer's velocity has a component along the source-to-observer line (zero if the observer is exactly abreast of the train, but nonzero when the observer is upstream or downstream). The shift is therefore nonzero at most points along the trajectory and only vanishes at the single instant when the observer is exactly perpendicular to the source. The "perpendicular at closest approach" version of no-Doppler is an instantaneous condition, not a sustained one.

We have now derived every formula needed for the first three worked examples. The fourth covers oblique Doppler, which uses the same source-moving formula with the source velocity component replaced by $u_s\cos\theta$.

BINDER Example 19

A train sounds a whistle of frequency 400Hz, approaching a railway platform at 10m/s. Take the speed of sound in still air as 340m/s. Find the frequency a stationary platform observer hears when the train is:

(a) approaching,

(b) receding.

Solution

(a) Apply the source-approaching formula with $u_s = 10\text{m/s}$, $f = 400\text{Hz}$, and $v = 340\text{m/s}$:

$$f' = f \times \frac{v}{v - u_s} = 400\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} - 10\text{m/s}} = 412.1\text{Hz}$$

(b) Apply the source-receding formula with the same numbers, switching the sign in the denominator:

$$f' = f \times \frac{v}{v + u_s} = 400\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} + 10\text{m/s}} = 388.6\text{Hz}$$

Dig deeper: The total swing in pitch as the train passes, from the approaching note to the receding note, is

$$412.1\text{Hz} - 388.6\text{Hz} = 23.5\text{Hz}$$

a clearly audible step that the platform observer would describe as the whistle "dropping its tone" the instant the train passes.

Making sense of the answer: *The shifts are roughly ± 3 per cent of the source frequency, the size of a small musical interval. A 23.5Hz swing on a 400Hz tone is about one semitone, just under a quarter-tone, well inside the resolution of any human ear. The asymmetry between the approaching and receding values is a hint that source motion is not purely a relative-velocity effect: $412.1Hz - 400Hz = 12.1Hz$ upwards on approach is not quite the same as $400Hz - 388.6Hz = 11.4Hz$ downwards on recede. The two differ at second order in u_s/v, an asymmetry we return to in REAL Example 21.*

Think like a physicist: *Whenever the source speed u_s is small compared with the wave speed v, both formulas reduce to the same first-order approximation $f' \approx f(1 \pm u_s/v)$. The student should learn to recognise this regime: at speeds well below the speed of sound (which is to say almost every speed encountered in everyday life), Doppler shifts are linear in source speed and approximately symmetric in sign. The non-linear regime, where the approach and recede shifts differ noticeably, sets in only when source speed becomes a substantial fraction of the wave speed, and is the regime where shock waves and sonic booms eventually live.*

REAL Example 20

A daladala on the Nyerere road sounds its horn, of true frequency 600Hz, while travelling at 25m/s along a straight section of road. A pedestrian standing on the grass verge hears the horn as the daladala approaches her and again as it recedes. Take the speed of sound in still air as 340m/s.

(a) Compute the apparent frequency the pedestrian hears before the daladala passes her, and the apparent frequency she hears after it has passed.

(b) Find the swing in pitch, and express it as a percentage of the true frequency of the horn.

(c) A passenger sitting inside the daladala hears the horn the whole time. Explain to the passenger, using the principle that governs Doppler effect, why she hears the horn at its true frequency of 600Hz at every moment, regardless of how fast the daladala is travelling.

Solution

(a) Apply the source-moving formulas to a stationary observer.

Before the daladala passes the pedestrian, the source approaches her at $u_s = 25\text{m/s}$:

$$f' = f \times \frac{v}{v - u_s} = 600\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} - 25\text{m/s}} = 647.6\text{Hz}$$

After the daladala passes, the source recedes from her at the same $u_s = 25\text{m/s}$:

$$f' = f \times \frac{v}{v + u_s} = 600\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} + 25\text{m/s}} = 558.9\text{Hz}$$

(b) The swing is the difference between the two apparent frequencies:

$$\Delta f = 648\text{Hz} - 559\text{Hz} = 89\text{Hz}$$

As a percentage of the true source frequency,

$$\frac{\Delta f}{f} = \frac{89\text{Hz}}{600\text{Hz}} = 0.148 = 14.8\%$$

The pedestrian hears the horn drop by almost fifteen per cent of its source frequency the instant the daladala passes her, the largest single-event pitch change a pedestrian standing on a Tanzanian road verge ever hears.

(c) The passenger inside the daladala shares the daladala's motion through the air. The horn, mounted on the daladala, also shares that motion. The horn and the passenger are therefore in the same rest frame: in their shared frame, neither is moving. The principle states that the apparent frequency is the wave speed seen by the observer divided by the wavelength as set by the source, both measured in the rest frame of the medium. In the passenger's frame, source and observer are at rest relative to each other, the air is what moves (backwards at 25m/s relative to the daladala), and the formal sub-case, "source and observer at the same velocity through the medium" forms no-Doppler effect scenario. The result is:

$$f' = f \times \frac{v + u_o}{v + u_s} = f \times \frac{v - u_s}{v - u_s} = f$$

(where the same speed has been substituted for u_s and u_o, with the same sign, because source and observer move identically through the air). The passenger hears 600Hz, not 648Hz or 559Hz, regardless of how fast the daladala travels. The Doppler effect requires that source and observer move **differently** through the medium; identical motion through the medium produces no shift.

Making sense of the answer: *Parts (a) and (b) give the largest pitch change in any of this section's worked examples, because the source speed of 25m/s is a sizeable seven per cent of the speed of sound. The asymmetry between the approach and recede shifts is now visible:* $648Hz - 600Hz = 48Hz$ *upward versus* $600Hz - 559Hz = 41Hz$ *downward. Part (c) is the conceptual punchline. A common student instinct is to think the passenger should hear a Doppler shift because the daladala is moving, but the moving daladala only matters relative to a stationary observer in the air outside; relative to the passenger, the horn is stationary. The Doppler effect is about motion through the medium, not about motion itself.*

Think like a physicist: *Two motions that produce the same relative velocity between source and observer can produce different Doppler shifts. The Doppler effect is asymmetric in source motion versus observer motion for a wave that travels in a medium; symmetry is only restored for a wave that travels in vacuum, namely electromagnetic radiation, which we treat in the next section. The asymmetry is a fingerprint of the medium: any wave that propagates in a medium carries this fingerprint, and any wave that propagates in vacuum does not.*

REAL Example 21

A Form Five student is cycling along the bus turnaround at Jangwani Secondary School at a steady 5m/s. Parked at the far end of the turnaround is a stationary Land Cruiser ambulance whose siren is set to a steady test tone of 800Hz for an electrical check. The student cycles directly towards the parked ambulance, passes it without slowing, and continues directly away from it on the other side. Take the speed of sound in still air as 340m/s.

(a) Compute the apparent frequency the cyclist hears while approaching the ambulance.

(b) Compute the apparent frequency she hears while receding from the ambulance.

(c) Compute, in two ways, the apparent frequency a stationary observer on the school field would hear if instead the ambulance were **moving** towards the field at 5m/s while the observer stood still. Compare the two predictions and comment on whether they agree.

Solution

(a) The source (the parked ambulance) is at rest in the air. The observer (the cyclist) moves towards the source at $u_o = 5\text{m/s}$. Apply the observer-approaching formula:

$$f' = f \times \frac{v + u_o}{v} = 800\text{Hz} \times \frac{340\text{m/s} + 5\text{m/s}}{340\text{m/s}} = 811.76\text{Hz}$$

(b) The cyclist continues past the ambulance and recedes from it at the same speed. Apply the observer-receding formula:

$$f' = f \times \frac{v - u_o}{v} = 800\text{Hz} \times \frac{340\text{m/s} - 5\text{m/s}}{340\text{m/s}} = 788.24\text{Hz}$$

(c) Now reverse the geometry. The ambulance moves towards the observer on the field at $u_s = 5\text{m/s}$, while the observer stands still. Apply the source-approaching formula:

$$f'_{\text{source-moving}} = f \times \frac{v}{v - u_s} = 800\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} - 5\text{m/s}} = 811.94\text{Hz}$$

Compare with part (a):

$$f'_{\text{observer-moving}} = 811.76\text{Hz}$$
$$f'_{\text{source-moving}} = 811.94\text{Hz}$$

The two answers are in close agreement, **differing by about 0.2Hz**. The source-moving case gives a slightly larger upward shift than the observer-moving case, for the same relative speed of 5m/s. The percentage difference is:

$$\frac{0.2}{812} = 0.00022 = 0.022\%$$

So the two scenarios are **not** exactly equivalent.

Making sense of the answer: *The 0.2Hz discrepancy is a real physical fact, not a rounding error. The source-moving formula has v in the numerator and $v - u_s$ in the denominator; the observer-moving formula has $v + u_o$ in the numerator and v in the denominator. The ratio between them, holding $u_s = u_o = u$, is*

$$\frac{f'_{source\text{-}moving}}{f'_{observer\text{-}moving}} = \frac{v/(v-u)}{(v+u)/v} = \frac{v^2}{(v-u)(v+u)} = \frac{v^2}{v^2-u^2}$$

which exceeds 1 by approximately u^2/v^2 for $u \ll v$. With $u = 5m/s$ and $v = 340m/s$, the discrepancy is $(5/340)^2 \approx 2.2 \times 10^{-4}$, exactly the 0.022 per cent we found. The Doppler effect for sound is genuinely asymmetric in source motion versus observer motion, even at the same relative speed.

Think like a physicist: *When two formulas almost agree but differ at second order, the disagreement carries information. The information here is that sound has a preferred frame: the rest frame of the air. Source motion and observer motion are physically distinguishable, and a careful measurement of f' at known u can in principle tell the experimenter whether it is the source or the observer that is moving. For light propagating in vacuum, by contrast, no such preferred frame exists; the Doppler formula for light is symmetric in source-versus-observer.*

HOT Example 22

A train approaches a level crossing at 120km/h along an east-west track. At the instant of interest, the train is exactly 300m east of the level crossing, with its whistle blowing a steady 640Hz tone. An observer is standing 400m north of the level crossing on the access road. Take the speed of sound in still air as 340m/s.

Find the apparent frequency the observer hears at this instant.

Solution

The figure below shows the setup.

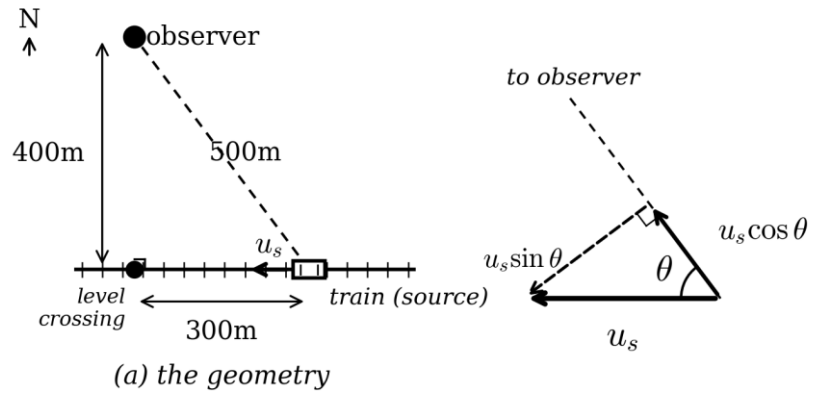


Figure: ***Geometry for HOT Example 22.*** *Panel (a): the train (source) is at 300m east of the level crossing; the observer is at 400m north of the level crossing; the source-to-observer line is the hypotenuse of a 3-4-5 right triangle of length 500m. Panel (b): the train's velocity u_s decomposes into a component $u_s cos\theta$ along the source-to-observer line (the only component that affects the Doppler shift) and a component $u_s sin\theta$ perpendicular to it (which does not).*

Convert the train's speed to SI units:

$$u_s = 120\text{km/h} = \frac{120000\text{m}}{3600\text{s}} = 33.33\text{m/s}$$

Find the source-to-observer distance:

From panel (a), the train, the level crossing, and the observer form a right triangle with legs 300m (east-west) and 400m (north-south). By Pythagoras, the hypotenuse is:

$$r = \sqrt{(300\text{m})^2 + (400\text{m})^2} = 500\text{m}$$

This is a 3-4-5 right triangle; the source-to-observer distance is 500m.

Find $cos\theta$:

The angle θ in panel (b) is the angle between the train's velocity (which points along the rails, westwards) and the source-to-observer line (which points from the train towards the observer). By inspecting the triangle in panel (a), θ is also the angle at the train inside the right triangle, between the east-west rail line and the hypotenuse. The cosine of that angle is the ratio of the adjacent leg (300m, along the rails) to the hypotenuse (500m, to the observer):

$$\cos\theta = \frac{\text{adjacent}}{\text{hypotenuse}} = \frac{300\text{m}}{500\text{m}} = 0.6$$

Find the component of velocity along the source-to-observer line:

Multiply the train's full speed by $\cos\theta$:

$$u_s\cos\theta = 33.33\text{m/s} \times 0.6 = 20.00\text{m/s}$$

This is the rate at which the train is closing the distance to the observer at this instant. The other component of the train's velocity, namely $u_s\sin\theta = 33.33\text{m/s} \times 0.8 = 26.66\text{m/s}$ (shown dashed in panel b), points sideways relative to the source-to-observer line. This perpendicular component carries the train across the line of sight rather than along it, so it does **not** contribute to the Doppler shift.

Apply the source-approaching Doppler formula:

Use $u_s\cos\theta$ in place of u_s:

$$f' = f \times \frac{v}{v - u_s\cos\theta} = 640\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} - 20.00\text{m/s}} = 680\text{Hz}$$

The observer hears the train's whistle at exactly 680Hz at this instant.

Making sense of the answer: *The component of the train's velocity along the source-to-observer line at this instant is 20m/s. That component, not the full 33.33m/s, is what enters the Doppler formula. The train is at 120km/h on the track, but only 72km/h "towards the observer" at this moment, because of the angle. The Doppler shift is therefore smaller than it would be if the observer stood directly on the track.*

Think like a physicist: *Doppler shift for an oblique source is governed by the velocity component along the source-to-observer line, not by the total speed. This generalises to a useful one-liner: the part of any motion that does not bring source and observer closer or further apart does not contribute to the Doppler shift. At the single instant when source and observer move exactly perpendicular to the line between them, the component along the line is zero and the apparent frequency equals the source frequency. This is the third of the three no-Doppler scenarios listed above, expressed in geometric language.*

This concludes the primary cases and their worked examples. The second half of this section turns to the two remaining complications, namely a wind blowing through the medium and a wave reflecting off a moving target.

Wind: the medium itself in motion

So far we have assumed the air stands still. Air rarely does. A sea breeze pushes inland from Bagamoyo every afternoon; the Maasai Mara highlands feed cool air down the Rift Valley by evening; the long-range monsoon flows north over Tanzania in summer and south in winter. Whenever the medium itself moves, the Doppler formula must be corrected, but the correction is small, intuitive, and easy to apply once stated properly.

The principle behind the correction is this. The wave speed v that appears throughout the still-air formula is the wave speed in the **rest frame of the medium**. When wind blows, the air is no longer at rest with respect to the ground, and the wave speed in the **rest frame of the ground**, where source and observer actually live and where their speeds are measured, is no longer v. The wind sweeps the air along with itself, carrying the wave with it, and the wave's ground speed in the source-to-observer direction picks up an extra contribution from the wind's velocity along that same direction.

Let **w** denote the component of the wind's velocity along the line from source to observer, signed so that $\mathbf{w > 0}$ means the wind blows *from the source toward the observer* and $\mathbf{w < 0}$ means the wind blows from the observer toward the source. The wave's ground speed in the source-to-observer direction is then:

$$\mathbf{V = v + w}$$

This V, not the still-air v, is what should appear in the Doppler formula when the experimenter chooses to measure the source and observer speeds in the ground frame (which is the natural choice). The source and observer velocities u_s and u_o continue to be measured as ground speeds, with the same sign convention as before. The Doppler formula with wind correction becomes:

$$\mathbf{f' = f \times \frac{V \pm u_o}{V \mp u_s}}, \qquad \text{where: } V = v + w$$

A wind of 30m/s in the source-to-observer direction lifts the effective wave speed from 330m/s to 360m/s. A wind of 30m/s from the observer toward the source drops the effective wave speed to 300m/s. A wind blowing perpendicular to the source-

to-observer line has zero component along that line, gives $w = 0$, and leaves the formula unchanged. Only the parallel component matters.

A subtle and counter-intuitive consequence of the wind correction is that a wind blowing in the same direction as the source's motion (a tailwind for the source) actually **reduces** the size of the Doppler shift. The reason is that the *wind makes wavefronts travel faster in the ground frame, so the source has less time to "catch up" with the wavefront it has just emitted, and the wavelength compression is therefore less severe.* Conversely, a headwind blowing back against the source's motion **increases** the Doppler shift, because the *wave moves more slowly in the ground frame and the source catches up with its own wavefronts more aggressively*. HOT Example 23 below exercises this subtlety with a worked numerical example at the Tabora railway station.

Round-trip Doppler: when the wave comes back

A different complication arises when the wave does not simply travel from source to observer but bounces off a reflector and returns. The reflector might be a cliff face (for the wave from a moving siren), a moving moth (for the chirp from a bat), a speeding car (for the radio pulse from a police speed gun), a moving submarine (for the pulse from a ship's sonar), or a moving red blood cell (for the ultrasound pulse from a medical Doppler probe). In each of these scenarios the original emitter eventually hears its own wave again, modified twice over.

The procedure is to apply the Doppler formula **twice in succession**, treating the reflector as a relay.

Step 1: *The reflector is the receiver of the wave coming from the original source.*

Apply the Doppler formula with the original source as the source and the reflector as the observer. Call the apparent frequency the reflector "hears" f_1.

Step 2. *The reflector immediately re-emits whatever wave it has received, at the same frequency f_1, in the reverse direction.*

The reflector is now the **source**, emitting f_1, and the original emitter (now playing the role of receiver) is the **observer**. Apply the Doppler formula again with the reflector as the source emitting f_1 and the original emitter as the observer. The frequency the original emitter ultimately hears is the round-trip Doppler-shifted frequency.

The reflector itself need not be stationary. If it is moving, then both Step 1 and Step 2 involve a moving observer (Step 1) or moving source (Step 2), and the formula must be applied with the reflector's velocity in mind for each step. The chain double-Doppler shift is approximately twice as large as a single-source Doppler shift for the same target velocity, a fact which the bat in HOT Example 24 will exploit to identify the velocity of its dinner with surprising precision, and which the police-speed-gun operator on the Morogoro Road in the applications section coming after the next one will exploit to issue the speeding ticket the chapter introduction warned the student about.

We are now equipped with all the tools required for the worked examples that follow.

HOT Example 23

A locomotive approaching the Tabora railway station sounds a horn of true frequency 660Hz, travelling at a ground speed of 30m/s towards a stationary passenger waiting on the platform. Take the speed of sound in still air as 330m/s.

(a) Compute the apparent frequency the passenger hears in still air.

(b) On a different day, a steady tailwind blows at 30m/s in the direction from the locomotive towards the platform, while the rest of the situation is unchanged. Compute the apparent frequency the passenger hears.

(c) On a third day, a steady crosswind of 30m/s blows perpendicular to the rail line, while the rest of the situation is again unchanged. Compute the apparent frequency the passenger hears, and explain in one sentence why.

Solution

(a) Still air, source approaching observer:

$$f' = f \times \frac{v}{v - u_s} = 660\text{Hz} \times \frac{330\text{m/s}}{330\text{m/s} - 30\text{m/s}} = 726\text{Hz}$$

The passenger hears the horn at 726Hz.

(b) The wind blows from the locomotive (the source) towards the platform (the observer), so the wind's component along the source-to-observer direction is $w = +30\text{m/s}$. The effective wave speed in the ground frame is:

$$V = v + w = 330\text{m/s} + 30\text{m/s} = 360\text{m/s}$$

Apply the wind-corrected Doppler formula, with the source speed and observer speed unchanged from part (a):

$$f' = f \times \frac{V}{V - u_s} = 660\text{Hz} \times \frac{360\text{m/s}}{360\text{m/s} - 30\text{m/s}} = 720\text{Hz}$$

The apparent frequency in the tailwind is 720Hz, only 6Hz **lower** than the still-air answer. The tailwind has **reduced** the Doppler shift from 66Hz above the source to 60Hz above.

(c) A crosswind blows perpendicular to the rail line, that is to say perpendicular to the source-to-observer direction. Its component along the source-to-observer direction is zero:

$$w = 0, \qquad V = v + 0 = 330\text{m/s}$$

The wind-corrected formula therefore reduces to the still-air formula, and the apparent frequency the passenger hears is:

$$f' = 726\text{Hz}$$

exactly the same as in part (a).

Only the wind's parallel component contributes to the Doppler correction; the perpendicular component has no effect.

Making sense of the answer: *Part (a) gives the baseline still-air Doppler shift, +66Hz upward. Part (b) gives the surprising result that a tailwind in the direction the source is moving reduces the shift, to +60Hz. The intuition is that the tailwind speeds up the wavefronts in the ground frame (they now propagate at 360m/s instead of 330m/s), so the source has less time to "catch up with" the wavefront it has just emitted, and the resulting wavelength compression is less severe. A headwind would do the opposite: the wave slows down in the ground frame, the source catches up more aggressively, and the Doppler shift grows. Part (c) confirms that only the parallel component of the wind matters; a perpendicular wind, however strong, has no effect on the apparent frequency.*

Think like a physicist: *The wind correction in this section is the simplest possible: the wind blows along or across the source-to-observer line, and the velocity decomposes cleanly into one parallel component (which contributes) and one perpendicular component (which does not).* ***For winds at intermediate angles, the same principle applies with cos θ: the contributing component of the wind is $w\cos\theta$, where θ is the angle between the wind direction and the source-to-observer line.*** *A wind blowing at 45 degrees to the rail line, at the same 30m/s magnitude, would contribute about $30\cos 45° = 21$m/s of effective ground-speed adjustment, giving an apparent frequency somewhere between the parallel-wind and zero-wind results.*

HOT Example 24

A bat hunting inside one of the limestone caves at Kondoa emits an ultrasonic chirp at a frequency of 40kHz. A moth, the bat's intended dinner, is flying through the cave at 5m/s. The cave air is still and at body temperature, with the speed of sound 340m/s.

(a) The moth is flying directly towards the bat at 5m/s. Find the frequency at which the bat receives its own chirp back after the chirp has reflected from the moth. State the chain of Doppler shifts that produces the answer.

(b) The bat now hovers in place and watches a second moth fly directly away from it at 5m/s. Find the frequency at which the bat receives the reflected chirp from this second moth. Compare with part (a) and explain how the bat tells the two moths apart.

Solution

(a) Apply the Doppler formula twice in succession, with the moth acting as a relay.

Step 1: *The bat is the source (stationary in the cave air), the moth is the observer (moving towards the bat at $u_m = 5$m/s).* The moth's velocity is along the source-to-observer line, so the simple observer-moving formula applies:

$$f_1 = f_0 \times \frac{v + u_m}{v} = 40000\text{Hz} \times \frac{340\text{m/s} + 5\text{m/s}}{340\text{m/s}} = 40588\text{Hz}$$

This is the frequency the moth, were it equipped with a bat-ear, would "hear" the chirp at.

Step 2: *The moth re-emits the chirp at f_1, this time playing the role of source moving towards the bat at the same $u_m = 5$m/s.* The bat is now the stationary observer. Apply the source-moving formula:

$$f_2 = f_1 \times \frac{v}{v - u_m} = 40588\text{Hz} \times \frac{340\text{m/s}}{340\text{m/s} - 5\text{m/s}} = 41194\text{Hz}$$

The bat receives the reflected chirp at 41194Hz, which is 1194Hz above the chirp it originally emitted.

Alternative solution

The two factors combine into a single expression for the *round-trip Doppler shift*,

$$\mathbf{f_2} = f_1 \times \frac{v}{v - u_m} = f_0 \times \frac{v + u_m}{v} \times \frac{v}{v - u_m} = \mathbf{f_0} \times \frac{\mathbf{v + u_m}}{\mathbf{v - u_m}} = 40000\text{Hz} \times \frac{345\text{m/s}}{335\text{m/s}} = 41194\text{Hz}$$

a one-line formula whose numerator and denominator both push the apparent frequency upwards (numerator by the moth-as-observer effect, denominator by the moth-as-source effect).

(b) The moth is now receding from the bat at $u_m = 5$m/s, the same speed in the opposite direction. Both halves of the round-trip formula flip sign:

$$f_2 = f_0 \times \frac{v - u_m}{v + u_m} = 40000\text{Hz} \times \frac{340\text{m/s} - 5\text{m/s}}{340\text{m/s} + 5\text{m/s}} = 38841\text{Hz}$$

The bat receives the reflected chirp at 38841Hz, which is 1159Hz **below** the chirp it originally emitted.

Comparing the two moths: the approaching moth shifts the returned chirp **up** by about 1.2kHz, and the receding moth shifts the returned chirp **down** by about 1.2kHz. The bat's auditory cortex, which is exquisitely sensitive to small frequency differences in the ultrasonic band, hears the two moths as two distinguishable tones: one above the 40kHz chirp it emitted, one below. The bat reads the **sign** of the shift to know which moth is coming towards it, and reads the **size** of the shift to know how fast.

Making sense of the answer: *The round-trip Doppler shift is approximately twice the single-Doppler shift for the same target speed, because the two factors in $(v + u_m)/(v - u_m)$ multiply together. For a moth at 5m/s, the round-trip shift is about 3 per cent of the carrier frequency; for a moth at 1m/s it is about 0.6 per cent. The bat's auditory system can resolve differences in pitch much finer than 0.1 per cent at ultrasonic frequencies, which is why bats can identify a hovering moth from a fast-darting moth in pitch darkness.*

Think like a physicist: *Every round-trip Doppler problem reduces to "apply the formula twice, with the reflector as a relay." The pattern repeats from acoustics (the bat in this example) to radar (the police speed gun in the applications section to come) to medical ultrasound (the blood-flow probe at Muhimbili). The mathematics is identical in each case; the wave is different. This is one of the best illustrations in classical physics of why the abstract formalism is worth learning: the formula derived for one situation transfers, intact and unmodified, to a dozen seemingly unrelated technologies.*

HOT Example 25

A police car chases a speeder along a straight section of the Morogoro Road towards a tall cliff face that rises perpendicular to the road. Both vehicles travel at 160km/h towards the cliff. The police-car siren emits a steady tone at 100Hz. A second car, driven by an off-duty inspector, follows the police car along the same straight section at 120km/h. Take the speed of sound in air as 330m/s.

Find:

(a) the frequency at which the inspector hears the police-car siren **directly**, by sound travelling from siren to inspector through the air,

(b) the frequency at which the inspector hears the **cliff-echo** of the same siren, by sound reflecting off the cliff face and travelling back to the inspector,

(c) the beat frequency between the two tones the inspector simultaneously hears.

Solution

Convert all speeds to SI units. The police car and the speeder both travel at:

$$u_p = 160\text{km/h} = \frac{160000\text{m}}{3600\text{s}} = 44.44\text{m/s}$$

The inspector's car travels at:

$$u_i = 120\text{km/h} = \frac{120000\text{m}}{3600\text{s}} = 33.33\text{m/s}$$

The geometry: all three vehicles move in the same direction along the road towards the cliff. The inspector's car is behind the police car (because the police car is the chaser, the speeder is ahead of the police, the inspector follows). Sound from the police-car siren travels backwards through the air to reach the inspector, and forwards through the air to reach the cliff.

(a) Direct sound from siren to inspector. The siren (source) is in front of the inspector (observer), and both move in the same direction. The police car, being faster, recedes from the inspector at the ground speed u_p. The inspector, moving in the same direction as the wave's source, approaches the source at the ground speed u_i.

Apply the general Doppler formula. Observer approaches the source: numerator carries a plus sign. Source recedes from the observer: denominator carries a plus sign.

$$f_{direct} = f \times \frac{v + u_i}{v + u_p} = 100\text{Hz} \times \frac{330\text{m/s} + 33.33\text{m/s}}{330\text{m/s} + 44.44\text{m/s}} = 97.03\text{Hz}$$

The inspector hears the direct siren at about 97Hz, that is to say slightly **below** the source frequency of 100Hz, because the source recedes faster than the observer pursues.

(b) The cliff echo arrives by reflection. Apply Doppler twice with the cliff as a stationary relay.

Step 1: The police car emits the siren and the cliff face is the stationary observer in the air.

Apply the source-approaching formula:

$$f_1 = f \times \frac{v}{v - u_p}$$

Step 2: The cliff re-emits the wave at f_1, acting as a stationary source. The inspector, now playing the role of observer. Apply the observer-approaching formula:

$$f_{echo} = f_1 \times \frac{v + u_i}{v} = f \times \frac{v}{v - u_p} \times \frac{v + u_i}{v} = f \times \frac{v + u_i}{v - u_p} = 100\text{Hz} \times \frac{330\text{m/s} + 33.33\text{m/s}}{330\text{m/s} - 44.44\text{m/s}} = 127.21\text{Hz}$$

The cliff-echo arrives at the inspector at 127.21Hz, well above the source frequency of 100Hz. The two factors (police approaching cliff, inspector approaching cliff) both push the echo upwards.

(c) The inspector simultaneously hears two pure tones: the direct siren at 97.0Hz and the cliff-echo at 127.2Hz. Two pure tones close in frequency superpose to produce beats. The beat frequency is the difference between the two tones:

$$f_{beat} = |f_{echo} - f_{direct}| = |127.21\text{Hz} - 97.03\text{Hz}| = 30.18\text{Hz}$$

The inspector hears, riding atop the average tone of $(97 + 127)/2 \approx 112$Hz, a slow throbbing at 30Hz, that is to say 30 loud-quiet cycles per second, exactly as Chapter 7 derived for the superposition of two close-frequency waves.

Making sense of the answer: *Three physical effects interact in this single problem. First, the direct sound is shifted downwards from 100Hz to 97Hz, because the source recedes from the observer faster than the observer chases the source. Second, the cliff-echo is shifted upwards from 100Hz to 127Hz, because both the original emission and the reflected arrival involve approach motions. Third, the two tones simultaneously incident on the inspector's ear produce a beat at the difference frequency of 30Hz. The Doppler effect, source reflection, and the principle of superposition (the last from Chapter 7) all play together in a single auditory experience: a continuous deep buzz at about 112Hz, throbbing thirty times a second. The inspector, if she is paying attention to the physics, can read the beat rate alone to deduce the closing speed of the police car towards the cliff.*

Think like a physicist: *The richest physical questions are the ones in which three or four pieces of physics that the student has learned separately suddenly turn out to be three or four pieces of the same phenomenon. This is one such question. The Doppler effect of this chapter, the round-trip relay of this section, the beat phenomenon from Chapter 7, all converge in the inspector's car on the Morogoro Road. Real-world acoustic environments are almost always multi-physics: the listener at the side of a busy highway hears direct engines, reflected engines, Doppler-shifted horns, beats between the Doppler-shifted tones, and ambient noise from many other sources, all simultaneously. Disentangling them mathematically is the whole point of having learned the physics. The exam question is a model of the world; the world is a richer version of the exam question.*

This section has done five things. It has derived Mr. Akilikubwa's blackboard equation from one principle, namely that the medium fixes the wave speed in its own rest frame and that source motion compresses wavelength while observer motion alters encounter rate. It has stated the three no-Doppler cases and shown that they all reduce to the same condition that source and observer share their velocity through the medium. It has corrected the formula for wind, that is to say for a medium that itself moves with respect to the ground. It has handled round-trip Doppler, where the wave bounces off a reflector and returns to be detected by the original emitter. And it has paid one of the four promises Mr. Akilikubwa made in the opening dialogue, by explaining how a bat in a Kondoa cave finds a moth in pitch darkness.

Three promises remain. Two of them, the colour of a distant galaxy and the speed of a car caught by a speed camera on the Morogoro Road, depend on Doppler shifts for **electromagnetic** waves rather than for sound. Electromagnetic waves, as the previous two sections established, propagate through vacuum at the fixed speed c and require no medium at all to do it. Without a medium there is no rest frame for the medium, and the asymmetry between source motion and observer motion that this section has just exploited disappears. The Doppler formula for light is symmetric in source and observer velocity, depending only on their relative velocity. The next section derives this symmetric formula, applies it to the redshift of distant galaxies (the first promise paid), introduces Hubble's law connecting redshift to distance, and treats the Doppler broadening of spectral lines. The final promise, about TBC radio at night, depends not on Doppler but on atmospheric propagation, and is treated in the section after that.

THE DOPPLER EFFECT FOR LIGHT, REDSHIFT, AND HUBBLE'S LAW

The previous section closed with the observation that the Doppler effect for light is fundamentally different from the Doppler effect for sound, because light does not require a medium in which to travel. The wave speed c is the same in every inertial frame; no preferred frame for the electromagnetic field has ever been found; and the asymmetry between source motion and observer motion, which we exploited in REAL Example 21 to expose the slight non-equivalence of approaching-source and approaching-observer scenarios at second order in u/v, disappears entirely for light. What is left is a Doppler formula that depends only on the **relative** velocity of source and observer, and that takes a particularly clean form for source-observer relative speeds much smaller than c.

This section does four things. It states the first-order Doppler formula for light, $\Delta\lambda/\lambda = v/c$, and reassures the student honestly that the rigorous relativistic formula reduces to this approximation for $v \ll c$ but requires special relativity to derive. It applies the formula to the redshift of a star, paying one of the four promises Mr. Akilikubwa made in the chapter's opening dialogue about the colour of a distant galaxy. It introduces Hubble's law, which links the recession velocity of a galaxy to its distance from us and gives the order-of-magnitude age of the universe almost as a free bonus. And it discusses the Doppler broadening of spectral lines, the phenomenon by which a hot gas of atoms in random thermal motion produces a smeared rather than a sharp emission line, and explains why the laboratory spectroscope at Mzumbe Secondary School cannot resolve atomic lines below a certain limit set entirely by temperature and atomic mass.

The first-order Doppler formula for light

When a source of light and an observer move with relative speed v along the line between them, with v measured as **positive for recession** and **negative for approach**, the wavelength of the light the observer measures is shifted from the wavelength the source emits according to:

$$\frac{\Delta\lambda}{\lambda} = \frac{v}{c},$$

where $\Delta\lambda = \lambda_{observed} - \lambda_{source}$ is the shift in wavelength and λ is the source wavelength. A **positive $\Delta\lambda$** means the *wavelength has stretched (the source recedes)*, and a **negative $\Delta\lambda$** means the *wavelength has compressed (the source approaches)*. Equivalently, the corresponding shift in frequency is

$$\frac{\Delta f}{f} = -\frac{v}{c},$$

with the opposite sign because frequency and wavelength are inversely related through $c = f\lambda$.

The formula is an approximation. The exact, rigorous formula that special relativity supplies is:

$$\frac{\lambda_{observed}}{\lambda_{source}} = \sqrt{\frac{1+\frac{v}{c}}{1-\frac{v}{c}}}$$

with v positive for recession. For $v \ll c$ this exact formula reduces to $1 + v/c$, which gives $\Delta\lambda/\lambda = v/c$, the approximation above. The error in the approximation is of order $(v/c)^2$, which for stars in our galaxy (typical speeds 100 to 1000km/s, so v/c of order 10^{-3}) is one part in a million, beyond the resolution of any school spectroscope. The approximation breaks down only when v becomes a substantial fraction of c, which happens for distant quasars and for cosmological redshifts above about 0.1. For everything in this section we use the first-order formula and flag, in the closing paragraph, where the approximation begins to bite.

Three differences between the light formula and the sound formula deserve emphasis.

1) Only relative velocity matters for light

There is no rest frame of the medium for an electromagnetic wave, because there is no medium. The Doppler shift for light depends only on the **relative** velocity of source and observer, not on either velocity separately. A source moving at speed v towards a stationary observer and a stationary source watched by an observer moving at speed v towards it produce **exactly** the same Doppler shift, to all orders in v/c. This is unlike the case of sound, where the air sets a preferred rest frame and source-vs-observer asymmetries appear at order $(u/v)^2$.

2) The formula is symmetric in v

Swap the roles of source and observer and the sign of v flips, but the formula keeps its form. This is a direct consequence of Einstein's principle that the laws of physics look the same in every inertial frame.

3) The Doppler shift for light has a real "limiting case" at $v = c$

At this limiting case, the exact formula gives $\Delta\lambda/\lambda \to \infty$ for recession and $\Delta\lambda/\lambda \to -1$ for approach. The first signals that a source receding at the speed of light would have its wavelength stretched without limit; the second, that an approaching source would have its wavelength contracted to zero. Neither case occurs in nature (no massive object can travel at c), but the closer one comes to c, the more dramatic the shifts become.

Redshift and blueshift: astronomy meets Doppler

The diagram below shows the same atomic spectral lines as they would appear from three different sources: a laboratory reference lamp at rest in the spectroscope, a star approaching the Earth at modest speed (the lines are shifted towards shorter wavelengths, hence the term **blueshift**), and a star receding from the Earth (the lines are shifted towards longer wavelengths, **redshift**). The dashed line traces the sodium D line at 5896Å in the laboratory reference and shows where the same physical line falls in each of the moving-source spectra.

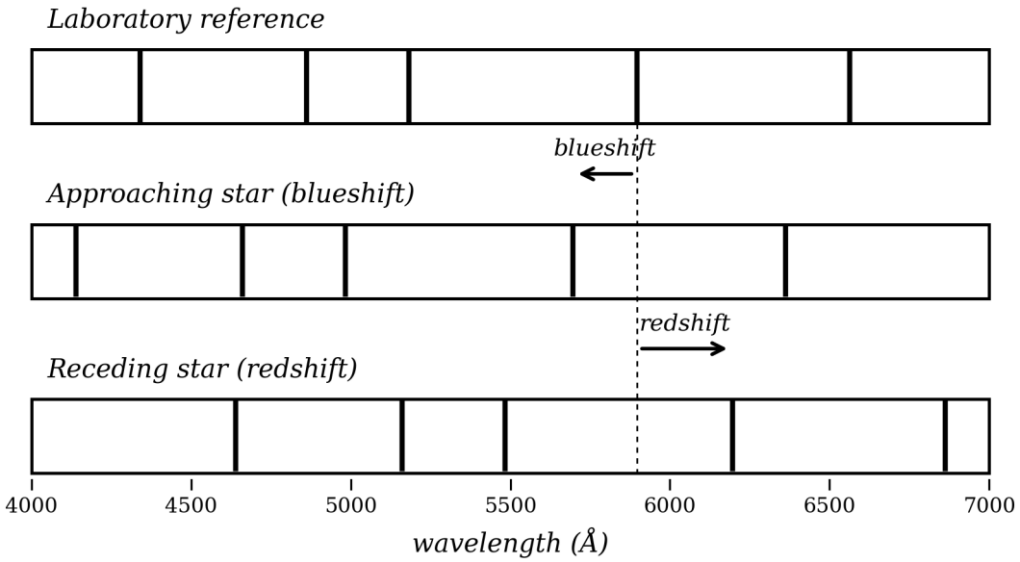


Figure: ***Spectral lines from three sources.*** *Top: laboratory reference lamp at rest. Middle: an approaching star, whose lines are blueshifted to shorter wavelengths. Bottom: a receding star, whose lines are redshifted to longer wavelengths. The dashed vertical line at 5896Å (the sodium D line in the laboratory) shows where the same line falls in each spectrum.*

The interpretation is straightforward. A star whose spectrum is blueshifted is approaching the Earth; a star whose spectrum is redshifted is receding. The size of the shift fixes the radial velocity through $v = c\,\Delta\lambda/\lambda$. The direction is read off the sign of $\Delta\lambda$.

There is one large fact that runs through every modern observation of the cosmos. *Almost every galaxy beyond our local group is redshifted, not blueshifted.* This means almost every galaxy is moving away from us. The few blueshifts that astronomers do measure correspond to nearby galaxies in our own gravitational neighbourhood (most famously the Andromeda galaxy, which is approaching us and will collide with the Milky Way in about four billion years). Beyond that nearby cluster, redshift dominates because the universe itself is expanding, a fact discovered by Edwin Hubble in 1929 and explained by the next subsection of this section.

The numerical worked example that follows is the textbook calculation of a stellar redshift.

BINDER Example 26

A spectral line from a distant star is identified as a sodium line. The same line, measured in a laboratory spectroscope at rest, has wavelength $\lambda_{lab} = 5896$ Å. When the spectroscope is pointed at the star, the same line is observed at $\lambda_{obs} = 6010$ Å. Take the speed of light to be $c = 3.0 \times 10^8$ m/s.

(a) Is the star approaching the Earth or receding from it?

(b) Compute the radial velocity of the star, that is to say the component of its velocity along the line of sight from the star to the Earth.

Solution

(a) Compute the shift:

$$\Delta\lambda = \lambda_{obs} - \lambda_{lab} = 6010\text{Å} - 5896\text{Å} = +114\text{Å}$$

The shift is positive: the observed wavelength is **longer** than the laboratory wavelength. This is a redshift, and the star is therefore **receding** from the Earth.

(b) Apply the first-order Doppler formula for light, $\Delta\lambda/\lambda = v/c$, solved for v:

$$v = c \times \frac{\Delta\lambda}{\lambda_{lab}} = 3.0 \times 10^8\text{m/s} \times \frac{114\text{Å}}{5896\text{Å}}$$

$$v = 3.0 \times 10^8\text{m/s} \times 0.0193 = 5.80 \times 10^6\text{m/s}$$

Converting to the unit astronomers prefer,

$$v = 5800\text{km/s}$$

The star is receding from the Earth at about 5800km/s (roughly 1.9 per cent of the speed of light).

This is fast in everyday terms (twenty thousand times the speed of a passenger jet) but still well below c, so the first-order formula is reliable here.

Making sense of the answer: *The sign of* $\Delta\lambda$ *tells the direction of motion: positive means recession, negative means approach. The size of* $\Delta\lambda$ *as a fraction of* λ *gives the speed as a fraction of* c*. The student should learn to read a spectroscope output in exactly this two-step way. The 1.9 per cent of* c *is a typical figure for a galaxy in the nearer half of the observable universe;*

nearby stars in our own galaxy are usually shifted by far less than this, by a fraction v/c *of order* 10^{-4}*, with corresponding shifts of fractions of an Ångström.*

Think like a physicist: *Notice that the radial velocity is what the Doppler shift gives, not the full velocity. A star could be moving transversely across the sky at any speed without any Doppler signature at all, just as the train in HOT Example 22 produced no Doppler shift from the perpendicular component of its velocity. To get the full velocity, an astronomer combines the Doppler-derived radial velocity with the transverse velocity (measured by tracking the star's apparent motion across the sky over many years) into a single three-dimensional velocity vector. This is the same vector-decomposition principle as the train at the level crossing, scaled up by a factor of ten quintillion.*

The first-order formula $\Delta\lambda/\lambda = v/c$ has served BINDER Example 26 well, because the star was receding at about 2 per cent of the speed of light, well inside the safe regime of the approximation. But astronomers regularly study sources whose recession velocity is a substantial fraction of c, and for such sources the approximation begins to break down noticeably. The worked example that follows takes a high-redshift quasar (a very luminous active galactic nucleus, named "quasi-stellar object" historically because the earliest examples looked like point-like stars in their photographs) and compares the first-order prediction against the exact relativistic Doppler formula. The discrepancy is large enough to matter for any quantitative work at that redshift.

HOT Example 27

A distant quasar in the constellation of Centaurus is observed with a sodium-D spectral line that is redshifted by the fraction $z = \Delta\lambda/\lambda_{lab} = 0.20$ from its laboratory wavelength.

(a) Use the first-order Doppler formula to estimate the recession velocity of the quasar.

(b) Use the exact relativistic Doppler formula to compute the exact recession velocity.

(c) Compare the two answers. By what percentage does the first-order formula overestimate the recession velocity?

Solution

(a) The first-order formula gives directly,

$$\frac{v}{c} \approx \frac{\Delta\lambda}{\lambda_{lab}} = z = 0.20$$

So:

$$v_{\text{first-order}} = 0.20 \times c = 0.20 \times 3.0 \times 10^8 \text{m/s} = 6.0 \times 10^7 \text{m/s} = 60000 \text{km/s}$$

The first-order formula predicts that the quasar is receding at 20 per cent of the speed of light.

(b) The exact relativistic Doppler formula:

$$\frac{\lambda_{obs}}{\lambda_{lab}} = \sqrt{\frac{1+\frac{v}{c}}{1-\frac{v}{c}}}$$

From the relativistic formula and the definition $z = \Delta\lambda/\lambda_{lab}$:

$$z = \frac{\lambda_{obs} - \lambda_{lab}}{\lambda_{lab}} = \frac{\lambda_{obs}}{\lambda_{lab}} - 1$$

$$\frac{\lambda_{obs}}{\lambda_{lab}} = z + 1 = 0.20 + 1 = 1.20$$

Let $\beta = v/c$.

Substitute into the exact formula and square both sides:

$$(1.20)^2 = \frac{1+\beta}{1-\beta}$$

$$\beta = 0.1803$$

The exact recession velocity is therefore:

$$v_{\text{exact}} = \beta \times c = 0.1803 \times 3.0 \times 10^8 \text{m/s} = 5.41 \times 10^7 \text{m/s} = 54100 \text{km/s}$$

The quasar is receding at 18.03 per cent of the speed of light, not the 20 per cent the first-order formula gave.

(c) The first-order formula gave 60000km/s; the exact formula gives 54100km/s. The first-order formula overestimates the true recession velocity by:

$$\Delta v = 60000 \text{km/s} - 54100 \text{km/s} = 5900 \text{km/s}$$

As a percentage of the exact value,

$$\frac{\Delta v}{v_{\text{exact}}} = \frac{5900}{54100} = 0.109 = 10.9\%$$

The first-order formula overestimates the recession velocity by about 11 per cent at this redshift. For accurate quasar kinematics, the exact relativistic formula must be used.

Making sense of the answer: *The size of the discrepancy is the practical statement of "where the approximation breaks down." A 10 per cent error is too large to ignore in any quantitative astronomical work. The error in the first-order formula scales roughly as* $z^2/2$*: at* $z = 0.01$ *the error is* 5×10^{-5} *(negligible), at* $z = 0.10$ *it is* 5×10^{-3} *(still tolerable for most*

work), at $z = 0.20$ it is 2×10^{-2} (the regime of this example, becoming uncomfortable), and at $z = 1$ the first-order formula fails completely. Modern cosmological surveys routinely deal with redshifts up to $z \approx 7$ for the most distant galaxies known, where the relativistic correction is dominant.

Think like a physicist: *The Doppler formula for light has the same general status as Newtonian mechanics: it is a low-velocity approximation to a more general theory that becomes essential at high velocity. The student should not despair that physical formulas have limited regimes of validity; this is the rule, not the exception. The skill the student must cultivate is recognising the regime of validity for each formula and switching to the more general formula when the regime breaks down. Newtonian mechanics is exact below 1 per cent of c; the first-order Doppler formula is good to 1 per cent below $z = 0.05$, to 10 per cent below $z = 0.20$, and useless above $z = 1$. Each formula carries its own boundary, and physics is the art of staying on the right side of the boundary or correctly stepping past it.*

Hubble's law

In 1929 the American astronomer Edwin Hubble compiled the recession velocities (measured by Doppler shift) and distances (estimated by various means, including the brightness of standard-candle stars called Cepheid variables) of many galaxies and noticed a striking pattern: *the recession velocity of a galaxy is approximately proportional to its distance from Earth.* Modern data, refined over nearly a century since Hubble's original work, confirm this relationship across hundreds of millions of parsecs and assemble it into what is now universally called **Hubble's law**:

$$\mathbf{v = H_0 d}$$

where v is the recession velocity of a galaxy, d is its distance from Earth, and H_0 is **Hubble's constant**, a fundamental cosmological parameter with current best-measured value approximately

$$H_0 \approx 70 \text{km/s per Mpc}$$

The unit "km/s per Mpc" reads as "every additional megaparsec (one Mpc = 3.086×10^{22} metres, or about 3.26 million light-years) of distance from us adds about 70km/s of recession velocity." In more student-friendly units, H_0 corresponds to about 22km/s per million light-years.

The figure below plots recession velocity against distance for a representative selection of galaxies. The straight line through the origin is Hubble's law. The data points scatter modestly around the line, in part because of measurement uncertainty and in part because of local gravitational motions superimposed on the cosmic expansion. The strong linear relationship dominates the scatter, especially at larger distances where the cosmic expansion overwhelms the local motions.

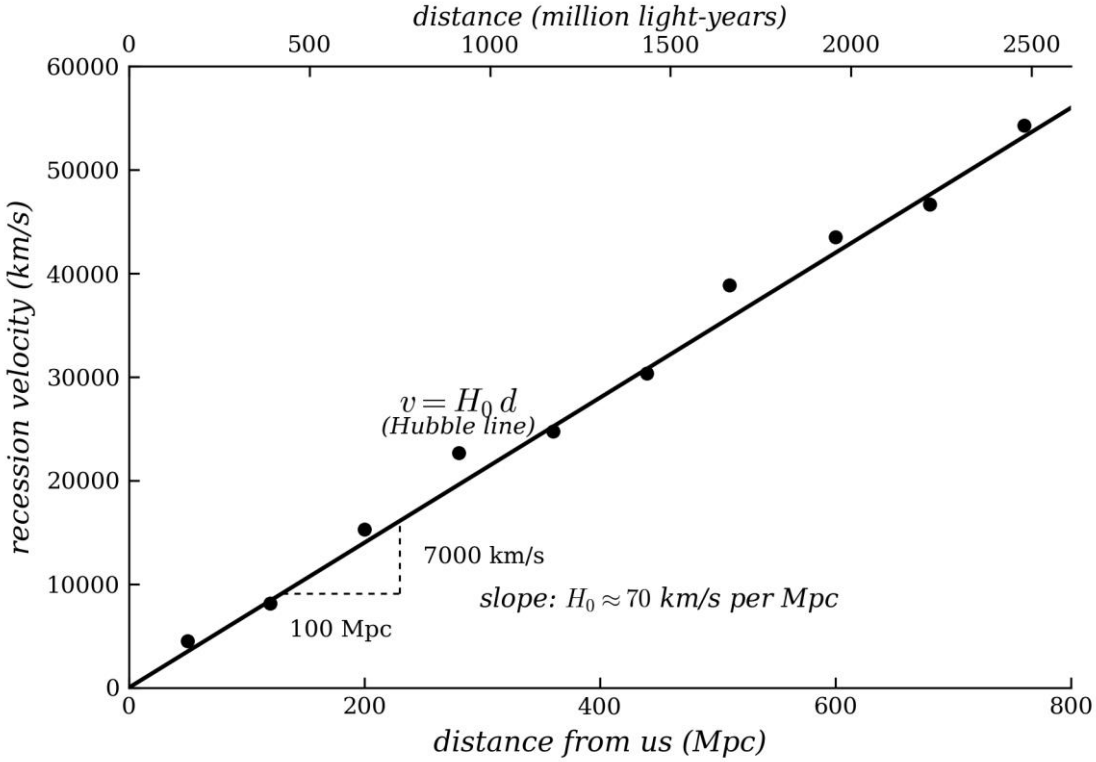


Figure: ***Recession velocity versus distance for a sample of galaxies.*** *The straight line is Hubble's law $v = H_0 d$ with $H_0 \approx 70$km/s per Mpc. The slope-triangle shows the meaning of the constant: every 100Mpc of additional distance corresponds to an additional 7000km/s of recession velocity. The upper axis converts distance from Mpc to millions of light-years using the conversion 1 Mpc ≈ 3.262 million light-years.*

The most striking property of Hubble's law, and the one that took half a century to understand, is what it says about the universe itself. If every galaxy is receding from us at a rate proportional to its distance, the natural conclusion is that the universe is **expanding uniformly**. This is not, as the student might first think, the consequence of us sitting at the centre of an explosion. Instead, it is the geometric consequence of space itself stretching out: every point in the universe sees every other point receding from it, with a recession speed proportional to the distance between them, exactly as Hubble's law describes. The student should picture not a bomb exploding outwards from a single centre, but a loaf of raisin bread rising in an oven: every raisin sees every other raisin receding, and the further away the receding raisin, the faster it recedes.

A useful by-product of Hubble's law is an estimate of the age of the universe. The quantity $1/H_0$, called the **Hubble time**, has units of time. Plugging in $H_0 \approx 70$km/s per Mpc:

$$\frac{1}{H_0} \approx \frac{1\text{Mpc}}{70\text{km/s}} = \frac{3.086 \times 10^{19}\text{km}}{70\text{km/s}} \approx 4.4 \times 10^{17}\text{s} \approx 14\text{billion years}$$

The Hubble time is the time it would take for the present rate of expansion to have produced the universe's present size, if the expansion has always proceeded at the same rate. The actual age of the universe, determined by more careful analysis of how the expansion rate has changed over cosmic history, comes out to be approximately 13.8 billion years, remarkably close to the Hubble-time estimate. The order of magnitude is set, and the precise correction requires detailed cosmological modelling that is well beyond the scope of this book.

The next worked example exercises Hubble's law on a concrete galaxy, the Virgo cluster of galaxies in the constellation of Virgo, the nearest large galaxy cluster to our own and a standard textbook target.

REAL Example 28

The Virgo cluster of galaxies sits at a measured distance of $d = 16.5$Mpc from the Milky Way. A spectral line of hydrogen (Hα, $\lambda_{lab} = 6563$Å) emitted by a galaxy in the Virgo cluster is observed on Earth at $\lambda_{obs} = 6587$Å. Take $H_0 = 70$km/s per Mpc and $c = 3.0 \times 10^8$m/s.

(a) Use Hubble's law to predict the recession velocity of the Virgo cluster.

(b) Compute the observed recession velocity directly from the Doppler shift of the Hα line.

(c) Compare the two answers. Comment on the size of the discrepancy.

Solution

(a) Hubble's law predicts $v = H_0\, d$. Substituting:

$$v_{predicted} = 70\text{km/s per Mpc} \times 16.5\text{Mpc} = 1155\text{km/s}$$

If the universe expanded uniformly with no other effects superimposed, the Virgo cluster would be receding from us at 1155km/s.

(b) Compute the wavelength shift:

$$\Delta\lambda = \lambda_{obs} - \lambda_{lab} = 6587\text{Å} - 6563\text{Å} = +24\text{Å}$$

Apply the first-order Doppler formula:

$$v_{observed} = c \times \frac{\Delta\lambda}{\lambda_{lab}} = 3.0 \times 10^8\text{m/s} \times \frac{24\text{Å}}{6563\text{Å}} = 1.097 \times 10^6\text{m/s} = 1097\text{km/s}$$

(c) The two values are:

$$v_{predicted} = 1155\text{km/s}, \qquad v_{observed} = 1097\text{km/s}$$

They differ by $\Delta v = 1155\ \text{km/s} - 1097\ \text{km/s} = 58$km/s, which is about 5 per cent of either value. The Hubble-law prediction is slightly higher than the actual observed recession.

Making sense of the answer: *The 5 per cent discrepancy is real, and it is the signature of what cosmologists call a* ***peculiar velocity****. The Virgo cluster does not move only with the smooth cosmic expansion; it also has a small additional motion through space, driven by the local gravitational pull of nearby mass concentrations. For Virgo specifically, this peculiar velocity points partly back towards us, slowing the cluster's net recession below the pure-Hubble value. Hubble's law is therefore strictly accurate only on average over many galaxies and over distances large enough that local gravitational effects average out. A single galaxy's redshift can deviate from the Hubble prediction by hundreds of km/s in either direction, but the trend across hundreds of galaxies is the clean linear law shown in the figure above.*

Think like a physicist: *The clean linear law sits underneath a noisy reality. Every individual galaxy carries a peculiar velocity on top of its cosmic-expansion velocity, and the peculiar velocity can point in any direction. Astronomers extract* H_0 *from large samples of galaxies, fitting a straight line through the cloud of data points, not from any single galaxy in isolation. This is the same principle that drives every physical-law measurement: the law is an idealisation, the data are noisy, and the constant is what comes out of the linear regression. The student should always ask, on encountering a physical law, "what is the noise that this law hides on top of?", and seek out at least one source of that noise. For Hubble's law, the noise is peculiar velocities. For Ohm's law, it is the temperature dependence of resistance. For the ideal-gas law, it is intermolecular forces. Every clean physical law has its noise; the law is the average, not the individual case.*

Doppler broadening of spectral lines

The third application of the light Doppler formula is to a single body of gas, hot enough to emit a spectral line, but with the atoms inside the gas moving with thermal velocities in random directions. Some atoms move towards the spectroscope and emit Doppler-blueshifted light; some move away and emit Doppler-redshifted light; most move at intermediate angles and produce intermediate shifts. The observed line is therefore not a single sharp wavelength but a **smeared** distribution centred on the laboratory wavelength, with a width controlled by the typical atomic speed.

The kinetic theory of gases (Chapter 4) gives the root-mean-square speed of atoms in an ideal gas at absolute temperature T as:

$$v_{rms} = \sqrt{\frac{3RT}{M}},$$

where $R = 8.314\text{J/(mol·K)}$ is the gas constant and M is the molar mass of the gas in kg/mol. Atoms move at this typical speed in random directions: about half of them have a component of velocity towards the spectroscope of order v_{rms}, and about half a component away of order v_{rms}. The observed spectral line is broadened by the combined Doppler shift, of magnitude approximately:

$$\Delta\lambda \approx \frac{2\lambda\, v_{rms}}{c}$$

This is the **Doppler width** of the spectral line. It depends only on three things: the wavelength of the line itself (λ), the absolute temperature of the gas (T), and the molar mass of the emitting atom (M). For a given line, lighter atoms move faster at any given temperature, and hotter gases move faster at any given mass; both effects broaden the line.

The next worked example exercises this formula. The example illustrates an important and counter-intuitive point: the **colder** gas can produce **broader** spectral lines, provided the atoms inside it are light enough.

HOT Example 29

Two laboratory discharge lamps emit the same green spectral line at $\lambda = 5461\text{Å}$. The first lamp contains mercury vapour at 200°C, the second krypton gas at 0°C. The molar masses are $M_{Hg} = 200\text{g/mol}$ for mercury and $M_{Kr} = 84\text{g/mol}$ for krypton. The gas constant is $R = 8.314\text{J/(mol·K)}$ and the speed of light is $c = 3.0 \times 10^8\text{m/s}$.

(a) Compute the root-mean-square atomic speed in each lamp.

(b) Compute the Doppler width of the spectral line in each lamp.

(c) Which lamp produces the sharper spectral line, and by what factor?

Solution

(a) Convert temperatures to absolute scale:

$$T_{Hg} = 200°\text{C} + 273 = 473\text{K}$$
$$T_{Kr} = 0°\text{C} + 273 = 273\text{K}$$

Convert molar masses to kg/mol:

$$M_{Hg} = 0.200\text{kg/mol}, \qquad M_{Kr} = 0.084\text{kg/mol}$$

Apply $v_{rms} = \sqrt{\frac{3RT}{M}}$ for each:

Mercury:

$$v_{rms,Hg} = \sqrt{\frac{3 \times 8.314\text{J/(mol·K)} \times 473\text{K}}{0.200\text{kg/mol}}} = 243\text{m/s}$$

Krypton:

$$v_{rms,Kr} = \sqrt{\frac{3 \times 8.314\text{J/(mol·K)} \times 273\text{K}}{0.084\text{kg/mol}}} = 285\text{m/s}$$

The krypton atoms move faster than the mercury atoms even though the mercury lamp is far hotter, because the krypton atoms are so much lighter.

(b) Apply the Doppler-broadening formula $\Delta\lambda = \frac{2\lambda v_{rms}}{c}$ for each. Convert $\lambda = 5461$ Å to metres: $\lambda = 5461 \times 10^{-10}\text{m} = 5.461 \times 10^{-7}\text{m}$.

Mercury:

$$\Delta\lambda_{Hg} = \frac{2 \times 5.461 \times 10^{-7}\text{m} \times 243\text{m/s}}{3.0 \times 10^8\text{m/s}} = 8.85 \times 10^{-13}\text{m} = 0.00885\text{Å}$$

Krypton:

$$\Delta\lambda_{\text{Kr}} = \frac{2 \times 5.461 \times 10^{-7}\text{m} \times 285\text{m/s}}{3.0 \times 10^{8}\text{m/s}} = 1.038 \times 10^{-12}\text{m} = 0.01038\text{Å}$$

(c) The sharper line has the smaller Doppler width. Comparing:

$$\frac{\Delta\lambda_{\text{Hg}}}{\Delta\lambda_{\text{Kr}}} = \frac{0.00885\text{Å}}{0.01038\text{Å}} = 0.853$$

The mercury lamp at 200°C produces the **sharper** line, narrower than the krypton lamp at 0°C by a factor of 0.853, equivalently the krypton line is broader by a factor of $1/0.853 = 1.17$.

Making sense of the answer: *The counter-intuitive result is that the cooler lamp (krypton at 0°C) produces broader spectral lines than the hotter lamp (mercury at 200°C), simply because krypton atoms are roughly 2.4 times lighter than mercury atoms. The line width scales as $\sqrt{T/M}$, and the krypton lamp wins on $1/M$ by a larger factor than it loses on T:*

$$\frac{\Delta\lambda_{\text{Hg}}}{\Delta\lambda_{\text{Kr}}} = \sqrt{\frac{\text{T}_{\text{Hg}}/\text{M}_{\text{Hg}}}{\text{T}_{\text{Kr}}/\text{M}_{\text{Kr}}}} = \sqrt{\frac{473/200}{273/84}} = \sqrt{\frac{2.365}{3.250}} = 0.853$$

confirming the result. The student should learn to recognise this trade-off: temperature and molar mass enter the line width through the single combination T/M, and a hot lamp made of heavy atoms can be just as quiet, spectroscopically speaking, as a cold lamp made of light atoms.

Think like a physicist: *Doppler broadening sets a fundamental limit on the resolution of any spectroscopic measurement performed on a hot gas. Higher-resolution spectroscopy at fixed wavelength requires cooling the gas (to slow the atoms) or selecting heavier atoms (so the random thermal motion is slower at any given temperature). The same principle drives modern atomic clocks: caesium atomic-clock laboratories cool their atoms to within a millionth of a kelvin of absolute zero so that Doppler broadening contributes negligibly to the precision of the frequency reference. Without cooling, the world's atomic clocks would lose seconds per day instead of seconds per million years.*

This section has paid the second of Mr. Akilikubwa's four chapter-opening promises by showing how the colour of a distant galaxy reveals its recession velocity through the redshift of its spectral lines, and how Hubble's law turns those redshifts into a panoramic picture of cosmic expansion. The third promise, the speed of a car caught by a speed camera on the Morogoro Road, depends on round-trip Doppler shifting of a radar pulse, the same physics the bat in Kondoa cave used to find its moth, but applied this time to electromagnetic waves and exploited by traffic police to issue speeding tickets. The fourth, the TBC radio at night, depends on atmospheric propagation, the way long-wavelength radio waves bend around the curvature of the Earth and bounce off the ionosphere, a quite different topic from Doppler that nevertheless rounds out the chapter. Both are brought together in the applications section that follows.

APPLICATIONS OF SOUND, EM WAVES, AND THE DOPPLER EFFECT

This final section of Chapter 9 introduces no new physics. Its purpose is to show the chapter's machinery, all of which the student has now derived and exercised, doing useful work in the real world. The seven applications below are arranged in the order the physics behind them appeared in this chapter, so a reader can quickly trace each one back to its theoretical home.

1. Ultrasound imaging in obstetrics

A pregnant woman attending her antenatal check at KCMC Hospital sees the inside of her own womb on a screen: the foetus, its heart beating, its head and limbs measurable to the millimetre. The image is constructed entirely from acoustic echoes. A piezoelectric transducer, held against the mother's abdomen, emits short pulses of ultrasound at a frequency between 2MHz and 15MHz, well above the audible range. The pulses travel into the body at the speed of sound in soft tissue, approximately 1540m/s, and partially reflect at every internal boundary where the acoustic impedance changes: skin-fat, fat-muscle, muscle-fluid, fluid-bone. Each reflected pulse returns to the transducer after a time proportional to twice the distance to the boundary that produced it. The machine measures this time, multiplies by half the wave speed, and assembles a two-dimensional image of the internal structures.

The choice of frequency in the megahertz range is a deliberate compromise. Higher frequencies give shorter wavelengths (at 15MHz the wavelength in tissue is about 0.1mm, smaller than the thickness of a human hair) and therefore sharper resolution, but they also attenuate more rapidly in tissue and so probe to shallower depths. Lower frequencies penetrate deeper but resolve less. Obstetric scans of a foetus deep inside the uterus typically use 3–5MHz; superficial scans of the thyroid or breast use 7–15MHz. The pedagogy connects to acoustic impedance from Chapter 9's opening section: only when impedance changes at a boundary does a reflection occur, and the brightness of the boundary on the screen is proportional to the impedance mismatch.

2. Sonar

The Tanzania People's Defence Force naval division operates sonar systems on its patrol boats off the coast at Dar es Salaam, mapping the seabed for navigation and detecting submerged objects. Civilian sonar serves the same physics in fishing trawlers off the Mafia archipelago, oceanographic survey vessels at the Pemba channel, and the harbour authorities at the Dar es Salaam port. A sonar transducer mounted on the underside of the vessel emits pulses of sound at frequencies between 20kHz

and 100kHz, depending on the application, into the water. The pulses travel through seawater at approximately 1500m/s, reflect off the seabed, fish schools, or any other object with a different acoustic impedance from the surrounding water, and return to the transducer.

The depth of a target is read off the round-trip time exactly as in obstetric ultrasound: a depth of 1km corresponds to a round-trip time of 2 × 1000m / 1500m/s ≈ 1.33s. Doppler sonar, an upgrade of the basic system, also measures the relative speed of any moving target the pulse bounces off, by reading the Doppler shift between the emitted and reflected frequencies. The chapter's round-trip Doppler formula, derived for the bat in the Kondoa cave, applies to the sonar without modification: only the medium and the wave are different.

3. Police radar speed gun

The road traffic police on the Morogoro Road operate radar speed guns, often hand-held devices held to the eye like a small camera. The device emits a continuous beam of microwave radiation at a fixed frequency (usually around 10GHz or 24GHz, in the so-called X-band or K-band of the radio spectrum) directly down the road, picks up the reflection from any oncoming or receding vehicle, and computes the vehicle's speed from the Doppler shift between the emitted and reflected frequencies. The round-trip relay used by the bat in HOT Example 24 of this chapter is exactly the principle in play.

The numerical scale is worth noting. At 24GHz, a car approaching at 100km/h produces a Doppler shift of,

$$\mathbf{\Delta f = \frac{2fv}{c}} = \frac{2 \times 24 \times 10^{9}\text{Hz} \times 27.8\text{m/s}}{3 \times 10^{8}\text{m/s}} \approx 4400\text{Hz}$$

a frequency difference well within the resolving power of any modern radio receiver. The speed gun reads this 4400Hz shift, divides by twice the carrier frequency, multiplies by the speed of light, and displays the answer in km/h to the operator in roughly a hundredth of a second. This is the third of Mr. Akilikubwa's chapter-opening promises paid: the speed of a car caught by a speed camera on the Morogoro Road is read directly from the same round-trip Doppler formula the bat used to find its moth.

4. Doppler weather radar

The Tanzania Meteorological Authority (TMA) operates weather radar at sites across the country, and the colour-coded weather maps shown on ITV nightly news bulletins are produced by these radars. The principle is the same as the police radar gun, but the targets are raindrops and ice particles inside clouds rather than vehicles on a road. A microwave pulse at around 5GHz (the C-band) is emitted from a ground-based dish, reflects off whatever raindrops or ice particles it meets, and returns to the receiver. The total power of the returned signal indicates how much water the radar beam encountered, which translates directly into rainfall intensity. The Doppler shift of the returned signal indicates the radial velocity of the rain or ice particles: positive shift means raindrops moving towards the radar, negative shift means moving away.

The colour-coding on the ITV weather forecast follows the velocity field: red colours indicate fast-approaching wind systems (rain bands closing in on the viewer), blue colours indicate fast-receding systems (rain moving away), and a rotational pattern of red against blue across a small region indicates a cyclonic vortex. Tropical cyclone tracking in the Indian Ocean basin, which affects the Tanzanian coast roughly twice per decade, relies critically on this Doppler velocity signature, since the rotational pattern reveals the cyclone's centre and the wind speeds at its eyewall.

5. Blood-flow ultrasound

A medical specialist at the Aga Khan Hospital who wishes to measure the blood flow through a patient's carotid artery applies an ultrasound transducer to the patient's neck, but adds a Doppler-shift readout to the basic imaging procedure described in the obstetric subsection above. The transducer emits a continuous pulse at typically 5MHz, and the red blood cells moving through the artery reflect a Doppler-shifted version of the pulse back to the transducer.

For a typical arterial blood flow of 1m/s and a 5MHz carrier, the round-trip Doppler shift is

$$\Delta f = \frac{2 f v}{v_{\text{tissue}}} = \frac{2 \times 5 \times 10^{6}\text{Hz} \times 1\text{m/s}}{1540\text{m/s}} \approx 6500\text{Hz},$$

a shift easily resolved by the receiver. The doctor reads off blood flow velocity in real time, can spot a partial blockage (the blood downstream of an obstruction speeds up to maintain volume flow, by the continuity principle from Chapter 10 of ***Foundations of Physics***), and can confirm the direction of flow (arterial flow into the head should always be in the same direction; a Doppler shift of the wrong sign is an immediate warning sign of vascular disease). The same physics, applied to the umbilical cord during pregnancy, monitors the blood supply between mother and foetus throughout pregnancy.

6. GPS as a Doppler-corrected positioning system

The Global Positioning System (GPS) used by every smartphone in Tanzania to find its location on a map depends critically on signals broadcast from a constellation of about 30 satellites orbiting the Earth at approximately 20200km altitude. Each satellite continuously broadcasts a microwave signal at 1575MHz (the so-called L1 frequency) carrying a precisely timed message identifying the satellite and the moment of transmission. A GPS receiver on the ground listens for these signals from at least four satellites simultaneously, computes the time delay between transmission and reception for each, and from those four delays solves for its own position in three dimensions plus the receiver clock offset.

Two corrections are essential, and without either the system would be useless. The first is the Doppler correction. Each satellite moves at about 3.9km/s in its orbit, with a radial component (the component of velocity along the line to the receiver) that can reach several hundred metres per second as the satellite rises or sets across the receiver's sky. At the 1575MHz L1 frequency, this radial motion produces a Doppler shift of a few kilohertz, well within the receiver's capability to measure. Without correcting for this shift, the receiver would misread the timing of the signal by an amount that translates, after the time-of-flight calculation, into a position error of several hundred metres at every measurement. Over the course of a day, uncorrected positions would drift by kilometres. The second correction is from the general theory of relativity whose discussion is beyond the scope of this book.

7. TBC radio at night: ionospheric skywave propagation

The fourth of Mr. Akilikubwa's chapter-opening promises was about TBC radio. A listener tuning a portable AM radio in Dar es Salaam at noon hears only the local Tanzanian stations broadcasting from transmitters within a few hundred kilometres. The same listener tuning the same radio at midnight often picks up stations broadcasting from Maputo, from Nairobi, from as far as Cape Town and Cairo. The same physical receiver, the same station antennas, but a very different reception. The reason has nothing to do with Doppler. It has to do with the ionosphere.

The Earth's upper atmosphere, between roughly 60km and 1000km altitude, is partially ionised by the Sun's ultraviolet and X-ray radiation. Free electrons in this ionised region (the **ionosphere**) act as a partial reflector for radio waves of the right frequency. The reflection is most efficient for AM-band frequencies (530–1700kHz, wavelengths of order a few hundred metres) bouncing off the so-called F-layer, the highest part of the ionosphere at about 250–400km altitude. During the day, the lower D-layer (around 60–90km altitude) absorbs these AM waves before they can reach the F-layer, and only line-of-sight propagation is possible: the listener hears only nearby transmitters. At night, the D-layer disappears (the Sun no longer ionises it, and the electrons recombine), the AM waves climb unimpeded to the F-layer, reflect, and return to the ground hundreds or thousands of kilometres from where they were transmitted. The listener in Dar es Salaam at midnight is receiving signals that have travelled by a single skip off the F-layer, sometimes by two skips with an intermediate ground reflection, from transmitters in southern Africa and beyond.

The physics is the same wave-propagation physics this chapter has used throughout: refraction (the ionosphere has a frequency-dependent refractive index that effectively bends the wave back down at appropriate angles), transmission (waves above a critical frequency penetrate the layer rather than reflecting), and reflection (waves below the critical frequency bounce back). Higher frequencies in the shortwave band (3–30MHz) reflect off lower layers and over shorter distances; lower frequencies penetrate higher. The student should not be surprised that this final phenomenon, paying the last promise of the opening dialogue, reuses only the wave physics already established earlier in the chapter; the chapter began with the observation that all waves obey one set of rules, and the chapter ends with a list of seven distinct technologies, each one a different combination of those same rules.

Mr. Akilikubwa's blackboard equation, the four promises he made in the chapter's opening dialogue, the eight sections that derived the physics, and the seven applications that completed it, now together form a single integrated treatment of sound waves, electromagnetic waves, and the Doppler effect for both. What remains is to see the machinery operating in full: mixed situations, hidden traps, numerical subtleties, sign-convention decisions, and real physical interpretations brought together in a final collection of miscellaneous worked examples.

MISCELLANEOUS WORKED EXAMPLES ON SOUND, ELECTROMAGNETIC WAVES, AND THE DOPPLER EFFECT

Example 30

(a) Explain why a person swimming below the surface of the sea at Coco Beach cannot hear a friend who is shouting from a boat in the air above, even though the swimmer's ears are working perfectly.

(b) An organ pipe has two consecutive resonant frequencies of 1372Hz and 1764Hz. Taking the speed of sound in air as 340m/s and neglecting end correction, determine whether the pipe is open or closed, calculate its fundamental frequency, and find its length.

Solution

(a) When a sound wave meets the boundary between two media, the fraction of its power that crosses depends on how closely the acoustic impedances $Z = \rho v$ of the two media match. Air has a very low acoustic impedance and water a very high one, the two differing by a factor of several thousand. Because the mismatch is so large, almost all the sound power arriving at the water surface from the air is reflected back into the air, and only about one part in a thousand is transmitted into the water. The small fraction that does enter the water spreads and weakens further, so by the time it reaches the submerged swimmer its intensity is far below the threshold of hearing.

(b) The difference of the two frequencies is:

$$\Delta f = 1764\text{Hz} - 1372\text{Hz} = 392\text{Hz}$$

Test for an open pipe:

An open pipe has consecutive resonances separated by the fundamental, so the fundamental frequency would be 392Hz. But $\frac{1372}{392} = 3.5$, which is not a whole number, so the pipe is not open.

Test for a closed pipe:

A closed pipe has consecutive resonances separated by $2f_1$, so $f_1 = \frac{392\text{Hz}}{2} = 196\text{Hz}$.

Then $\frac{1372}{196} = 7$ and $\frac{1764}{196} = 9$, both odd whole numbers. Hence the pipe is closed, sounding its seventh and ninth harmonics. The fundamental frequency is 196Hz.

For a closed pipe the fundamental is $f_1 = \frac{v}{4L}$, so the length is:

$$L = \frac{v}{4f_1} = \frac{340\text{m/s}}{4 \times 196\text{Hz}} = 0.434\text{m} = 43.4\text{cm}$$

Example 31

(a) Sir Isaac Newton's calculation of the speed of sound in air gave a value about sixteen per cent below the measured value. State the physical assumption Newton made that was wrong, identify the correction supplied by Laplace, and explain why the correction was necessary.

(b) An observer walks along the straight line joining two identical stationary sources, each emitting a steady note of frequency 500Hz. She hears 4.0 beats per second. Taking the speed of sound in still air as 340m/s, calculate her walking speed.

Solution

(a) Newton assumed that the compressions and rarefactions of a sound wave take place isothermally, that is, slowly enough for heat to flow and keep the gas at the temperature of its surroundings. This made the elastic property of the gas equal to its pressure P. Laplace recognised that a sound wave oscillates far too rapidly for any appreciable heat to flow during a single cycle, so each compression is in fact adiabatic, and the correct elastic property is the adiabatic bulk modulus γP rather than P. Replacing P by γP inside the square root multiplies the predicted speed by $\sqrt{\gamma}$, and for air $\sqrt{1.40} \approx 1.18$, which closes the sixteen-per-cent gap and brings the formula into agreement with experiment.

(b) As the observer walks along the line, she approaches one source and recedes from the other at the same speed u. For the source she approaches, the apparent frequency is raised:

$$f_1 = \frac{f(v + u)}{v}$$

For the source she recedes from, the apparent frequency is lowered:

$$f_2 = \frac{f(v - u)}{v}$$

The two notes are close in frequency and superpose to give beats at the difference frequency:

$$\mathbf{f_{beat}} = f_1 - f_2 = \frac{f(v + u)}{v} - \frac{f(v - u)}{v} = \mathbf{\frac{2fu}{v}}$$

Solve for u:

$$u = \frac{f_{\text{beat}} \times v}{2f} = \frac{4.0\text{Hz} \times 340\text{m/s}}{2 \times 500\text{Hz}} = 1.36\text{m/s}$$

The observer is walking at 1.36m/s.

Example 32

(a) Explain why a closed organ pipe sounds a note one octave lower than an open organ pipe of the same length, and explain why a closed pipe produces only the odd harmonics while an open pipe produces all of them.

(b) The speed of sound in a certain gas is measured as 353m/s at a temperature of 27°C. The molar mass of the gas is 28.0g/mol and the molar gas constant is $R = 8.314\text{J/(mol} \cdot \text{K)}$. Calculate the ratio of the principal specific heat capacities of the gas, and state, with a reason, whether the gas is monatomic or diatomic.

Solution

(a) A closed pipe has a displacement node at its closed end and a displacement antinode at its open end, so the shortest standing wave that fits has just one quarter of a wavelength inside the pipe, giving a fundamental frequency $f = \frac{v}{4L}$. An open pipe has a displacement antinode at each end, so its shortest standing wave has half a wavelength inside, giving a fundamental $f = \frac{v}{2L}$, exactly twice that of a closed pipe of the same length, which is one octave higher. The closed pipe produces only odd harmonics because every allowed mode must keep a node at the closed end and an antinode at the open end, and only odd multiples of the quarter-wavelength satisfy both conditions at once. An open pipe, with an antinode required at each end, admits every multiple of the half-wavelength, both odd and even, so it produces the complete set of harmonics.

(b) The speed of sound in an ideal gas is $v = \sqrt{\frac{\gamma RT}{M}}$. Solve for γ:

$$\gamma = \frac{Mv^2}{RT}$$

Convert the data: $M = 28.0\text{g/mol} = 0.0280\text{kg/mol}$ and $T = (27 + 273)\text{K} = 300\text{K}$.

$$\gamma = \frac{0.0280\text{kg/mol} \times (353\text{m/s})^2}{8.314\text{J/(mol}\cdot\text{K)} \times 300\text{K}} = 1.40$$

The ratio of the specific heat capacities is 1.40. A monatomic gas has $\gamma = 1.67$ and a diatomic gas has $\gamma = 1.40$, so the gas is **diatomic**.

Example 33

(a) Explain why the siren of an ambulance racing along United Nations Road sounds higher in pitch while the ambulance approaches a listener and lower in pitch after it has passed, and explain why the change in pitch is sudden rather than gradual.

(b) A single diesel generator produces a sound level of 88dB at a distance of 5.0m. Treating the generators as incoherent point sources, calculate the distance from a cluster of five such generators at which the sound level falls to 85dB.

Solution

(a) The siren emits wavefronts at a fixed rate. While the ambulance approaches the listener, it moves a little closer between one emission and the next, so the wavefronts ahead of it are crowded into a shorter distance; the listener meets shorter wavelengths, and since the speed of sound in the air is unchanged, a shorter wavelength means a higher frequency. After the ambulance has passed, it moves away between emissions, the wavefronts behind it are stretched, the wavelength is longer and the frequency lower. The change is sudden because, at the instant the ambulance passes the listener, the motion along the line joining them switches abruptly from approach to recession, so the apparent frequency jumps straight from its raised value to its lowered value with nothing in between.

(b) First combine the five generators. Five incoherent equal sources raise the level above that of one source by:

$$\Delta L = 10\log 5 = 10 \times 0.699 \approx 7.0\text{dB}$$

So the cluster produces $88\text{dB} + 7.0\text{dB} = 95\text{dB}$ at 5.0m.

The level must now fall from 95dB to 85dB, a drop of:

$$\Delta L = 85\text{dB} - 95\text{dB} = -10\text{dB}$$

For a point source the level changes with distance as $\Delta L = -20\log\left(\frac{r_2}{r_1}\right)$:

$$-10\text{dB} = -20\log\left(\frac{r_2}{5.0\text{m}}\right)$$

$$r_2 = 15.8\text{m}$$

A worker must stand at least 15.8m from the cluster of five generators for the sound level to fall to 85dB.

Example 34

(a) A workshop installs more identical machines. Explain why doubling the number of machines raises the measured sound level by only about 3 decibels, and explain why making the workshop sound half as loud to the ear cannot be achieved simply by removing half the machines.

(b) Calculate the speed of sound in a steel railway rail, for which the Young's modulus is $2.0 \times 10^{11}\text{Pa}$ and the density is 7800kg/m^3, and in seawater, for which the bulk modulus is $2.2 \times 10^9\text{Pa}$ and the density is 1030kg/m^3. Hence find the ratio of the two speeds.

Solution

(a) Incoherent sources add their intensities, so doubling the number of machines doubles the total intensity. On the decibel scale a doubling of intensity raises the level by $10\log 2 \approx 3\text{dB}$, regardless of how many machines there were to begin with. The decibel scale is logarithmic and roughly matches the response of the ear, but the ear perceives a doubling of loudness only for a level increase of about 10dB, which corresponds to a tenfold increase in intensity. To make the workshop sound only half as loud, the intensity must be cut to about one tenth, which means removing about nine machines out of every ten. Removing only half the machines lowers the level by just 3dB, far short of the 10dB the ear needs to register a halving of loudness.

(b) In a solid the elastic property is the Young's modulus and the wave speed is $v = \sqrt{\frac{Y}{\rho}}$. For the steel rail:

$$v_{\text{steel}} = \sqrt{\frac{Y}{\rho}} = \sqrt{\frac{2.0 \times 10^{11}\text{Pa}}{7800\text{kg/m}^3}} = 5064\text{m/s}$$

In a liquid the elastic property is the bulk modulus and the wave speed is $v = \sqrt{\frac{K}{\rho}}$. For seawater:

$$v_{water} = \sqrt{\frac{K}{\rho}} = \sqrt{\frac{2.2 \times 10^9 Pa}{1030 kg/m^3}} = 1461 m/s$$

The ratio of the two speeds is:

$$\frac{v_{steel}}{v_{water}} = \frac{5064 m/s}{1461 m/s} = 3.47$$

Sound travels about 3.5 times faster in the steel rail than in seawater.

Example 35

(a) Radio waves, visible light, X-rays and gamma rays are all members of the electromagnetic family. Explain why all of them travel at the same speed in a vacuum, and yet behave very differently when they meet matter.

(b) An open pipe of length 36.0cm and a closed pipe of length 27.5cm, both of the same diameter, are sounded so that the first overtone of the open pipe is in unison with the first overtone of the closed pipe. Calculate the end correction at one open end, and hence the internal radius of the pipes.

Solution

(a) Every electromagnetic wave in a vacuum travels at the speed $c = \frac{1}{\sqrt{\mu_0 \varepsilon_0}}$, where μ_0 and ε_0 are properties of empty space alone. These two constants do not depend on the frequency or wavelength of the wave, so every member of the family, from the longest radio wave to the shortest gamma ray, travels at exactly the same speed in a vacuum; the members differ only in wavelength and frequency. How a wave interacts with matter, however, depends strongly on its wavelength compared with the size of the objects and structures it meets, and on the energy each oscillation can deliver. A long radio wave diffracts around hills and buildings, visible light is absorbed and reflected by ordinary surfaces, and a short X-ray passes through flesh but is stopped by bone. Same speed in a vacuum, different wavelength, different behaviour in matter.

(b) Let e be the end correction at one open end, so $e = 0.6r$. The effective lengths are:

$$L_{eff}(open) = L_o + 2e = 0.360m + 2e$$
$$L_{eff}(closed) = L_c + e = 0.275m + e$$

The first overtone of the open pipe is its second harmonic, of frequency $\frac{2v}{2L_{eff}(open)} = \frac{v}{L_{eff}(open)}$.

The first overtone of the closed pipe is its third harmonic, of frequency $\frac{3v}{4L_{eff}(closed)}$.

Setting the two equal for unison:

$$\frac{v}{L_{eff}(open)} = \frac{3v}{4L_{eff}(closed)}$$
$$4L_{eff}(closed) = 3L_{eff}(open)$$
$$4(L_c + e) = 3(L_o + 2e)$$
$$2e = 4L_c - 3L_o$$
$$e = \frac{4L_c - 3L_o}{2} = \frac{4 \times 0.275m - 3 \times 0.360m}{2} = 0.010m$$

The end correction at one open end is 0.010m, or 1.0cm.

The radius follows from $e = 0.6r$:

$$r = \frac{e}{0.6} = \frac{0.010m}{0.6} = 0.0167m = 1.67cm$$

Example 36

(a) Explain why an astronaut standing on the surface of the Moon can see a colleague strike two rocks together but cannot hear the strike, and use the same physics to explain why sound is classified as a mechanical wave while light is not.

(b) A spectral line whose laboratory wavelength is 5000Å is observed in the light of a star belonging to a binary system to oscillate between 4999.6Å and 5000.4Å over one orbital period. Taking the speed of light as $3.0 \times 10^8 m/s$, calculate the orbital speed of the star.

Solution

(a) Sound is a mechanical wave, which means it is a disturbance of a material medium and cannot exist without one. Light is an electromagnetic wave, an oscillation of electric and magnetic fields that needs no medium at all. The Moon has no atmosphere, so on its surface there is no material medium for a sound wave to disturb, and the strike of the rocks produces no sound that can travel. The light reflected from the rocks, needing no medium, crosses the vacuum freely, so the colleague's action is seen but not heard. The same contrast fixes the classification: a wave that requires a medium to carry it is mechanical, and sound, which dies the instant its medium is removed, belongs to that class, while light, which crosses empty space unaided, does not.

(b) The line is most blueshifted when the star moves directly towards the Earth and most redshifted when it moves directly away, both at the full orbital speed v. The shift in each direction has magnitude $\Delta\lambda = \frac{v\lambda}{c}$, so the full swing from the shortest to the longest observed wavelength is:

$$\lambda_{max} - \lambda_{min} = 2\Delta\lambda = \frac{2v\lambda}{c}$$

Solve for v:

$$v = \frac{c(\lambda_{max} - \lambda_{min})}{2\lambda}$$

$$v = \frac{3.0 \times 10^8 m/s \times (5000.4Å - 4999.6Å)}{2 \times 5000Å} = 2.4 \times 10^4 m/s$$

The star moves in its orbit at $2.4 \times 10^4 m/s$, or 24km/s.

Example 37

(a) Motion does not always produce a Doppler shift. State and explain the three situations in which a source, an observer, or both, are in motion and yet no change of frequency is heard.

(b) An air-traffic control radar at Julius Nyerere International Airport transmits microwave pulses at a frequency of 9.0GHz towards an aircraft 30km away. Calculate the wavelength of the pulses and the time that elapses between the transmission of a pulse and the return of its reflection.

Solution

(a) There are three such situations.

1) When both source and observer are at rest in the medium, the wavefronts are emitted and met at the same rate, and the frequency is unchanged.
2) When source and observer move with the same velocity through the medium, they are at rest relative to each other; a passenger seated in a daladala that is sounding its horn hears the horn at its true frequency however fast the daladala travels, because the horn and the passenger share the same motion through the air.
3) When the velocity of the source or observer is directed exactly at right angles to the line joining the two, as when a source moves along a circle with the observer at its centre, the distance between them is momentarily not changing, the component of velocity along that line is zero, and no shift is produced.

(b) Treat the air as effectively a vacuum, so the wave speed is $c = 3.0 \times 10^8 m/s$.
Wavelength:

$$\lambda = \frac{c}{f} = \frac{3.0 \times 10^8 m/s}{9.0 \times 10^9 Hz} = 0.033m = 3.3cm$$

The pulse travels to the aircraft and back, a total distance of $2 \times 30km = 6.0 \times 10^4 m$, at the speed of light:

$$t = \frac{2d}{c} = \frac{6.0 \times 10^4 m}{3.0 \times 10^8 m/s} = 2.0 \times 10^{-4} s = 0.20ms$$

The pulses have a wavelength of 3.3cm, and each reflection returns 0.20 milliseconds after the pulse is sent.

Example 38

(a) Explain why a listener with a portable radio in Dar es Salaam can pick up distant stations from Nairobi or Maputo at midnight but hears only local stations at noon.

(b) On a cold morning at 13°C, two organ pipes of slightly different lengths produce 6.0 beats per second when sounded together. Calculate the number of beats per second the same two pipes produce on a hot afternoon when the air temperature has risen to 38°C.

Solution

(a) Radio waves of the medium-wave broadcasting band can travel far beyond the horizon only by reflecting off the ionosphere, the electrically charged upper layers of the atmosphere. During the day the Sun's radiation creates a low absorbing layer in the ionosphere that soaks up these radio waves before they can reach the higher reflecting layer, so the listener hears only nearby stations whose signals arrive directly. At night this low absorbing layer disappears, the radio waves rise unhindered to the high reflecting layer, and they are bounced back to the ground hundreds of kilometres away. The listener in Dar es Salaam at midnight is therefore receiving signals that have reflected off the ionosphere from transmitters in distant cities, a path that the daytime absorbing layer had blocked.

(b) The frequency of an organ pipe is proportional to the speed of sound in the air inside it, and the speed of sound is proportional to the square root of the absolute temperature, $v \propto \sqrt{T}$.

When the air warms, the frequency of each pipe rises in the proportion $\sqrt{T}$, and since the beat frequency is the difference of the two pipe frequencies, it is multiplied by the same factor:

$$f_{beat}(T_2) = f_{beat}(T_1) \times \sqrt{\frac{T_2}{T_1}}$$

Convert the temperatures: $T_1 = (13 + 273)K = 286K$ and $T_2 = (38 + 273)K = 311K$.

$$f_{beat}(T_2) = 6.0\ \text{beats/s} \times \sqrt{\frac{311K}{286K}} = 6.3\ \text{beats/s}$$

The two pipes produce about 6.3 beats per second on the hot afternoon.

Example 39

(a) Explain why the speed of sound in air depends strongly on the temperature of the air but not at all on its pressure.
(b) A train sounds a whistle of true frequency 510Hz. As it travels along a straight track at constant speed, the frequency heard by a stationary observer while the train approaches is to the frequency heard after the train has passed in the ratio 9:8. Taking the speed of sound in still air as 340m/s, calculate the speed of the train and the two apparent frequencies.

Solution

(a) The speed of sound in an ideal gas can be written $v = \sqrt{\frac{\gamma RT}{M}}$, in which the density has cancelled out, leaving only the absolute temperature, the molar mass, and the constant γ. The speed therefore rises with temperature, in proportion to the square root of the absolute temperature, because warmer molecules move faster and pass the disturbance along more quickly. Pressure does not appear in this form at all. Although the alternative form $v = \sqrt{\frac{\gamma P}{\rho}}$ seems to show a dependence on pressure, at a fixed temperature the density of a gas is itself proportional to its pressure, so the ratio $\frac{P}{\rho}$ stays constant; doubling the pressure doubles the density, and the speed is unchanged. Temperature alters the speed of sound, pressure does not.

(b) For the approaching train the apparent frequency is $f' = \frac{fv}{v - u}$; for the receding train it is $f' = \frac{fv}{v + u}$. The ratio of the two is:

$$\frac{f'(\text{approach})}{f'(\text{recede})} = \frac{\frac{fv}{v - u}}{\frac{fv}{v + u}} = \frac{v + u}{v - u}$$

Set this equal to $\frac{9}{8}$:

$$\frac{v + u}{v - u} = \frac{9}{8}$$
$$8(v + u) = 9(v - u)$$
$$17u = v$$
$$u = \frac{v}{17} = \frac{340m/s}{17} = 20m/s$$

The train is travelling at 20m/s.
The two apparent frequencies are:

$$f'(\text{approach}) = \frac{fv}{v - u} = \frac{510Hz \times 340m/s}{340m/s - 20m/s} = 541.9Hz$$
$$f'(\text{recede}) = \frac{fv}{v + u} = \frac{510Hz \times 340m/s}{340m/s + 20m/s} = 481.7Hz$$

The observer hears 541.9Hz on approach and 481.7Hz on recession.

Example 40

(a) Explain why an end correction must be applied to the open end of a pipe, and explain why a student measuring the speed of sound with a resonance tube is advised to use two consecutive resonance lengths rather than a single one.
(b) A radar speed gun on the Morogoro Road operates at a frequency of 24GHz. It detects a Doppler shift of 3.6kHz in the wave reflected from an approaching car. Taking the speed of light as $3.0 \times 10^8 m/s$, calculate the speed of the car.

Solution

(a) The displacement antinode at the open end of a pipe does not sit exactly at the geometric end. The air just inside the open end is free to move, and its motion extends a little way into the air beyond the pipe, so the antinode lies a small distance outside the geometric end. This extra distance is the end correction, and it must be added to the measured length to give the true effective length of the air column. A student measuring the speed of sound with a resonance tube cannot easily measure this end correction directly. If, however, two consecutive resonance lengths are recorded, the difference between them is exactly half a wavelength and the unknown end correction cancels out of the subtraction. Using two lengths rather than one therefore removes the systematic error the end correction would otherwise introduce.

(b) For a radar gun the wave is shifted twice, once when the moving car receives it and once when the car reflects it, so the detected shift for an approaching target is $\mathbf{\Delta f = \frac{2fu}{c}}$. Solve for u:

$$u = \frac{\Delta f \times c}{2f}$$

$$u = \frac{3.6 \times 10^3 \text{Hz} \times 3.0 \times 10^8 \text{m/s}}{2 \times 24 \times 10^9 \text{Hz}} = 22.5 \text{m/s} = 81 \text{km/h}$$

The car is travelling at 22.5m/s, or about 81km/h.

Example 41

(a) Explain why the spectral lines of almost every distant galaxy are observed to be shifted towards the red end of the spectrum, and state what an astronomer learns about a galaxy from the size and the direction of the shift.

(b) For nitrogen gas at 300K, with $\gamma = 1.40$, molar mass 28.0g/mol and $R = 8.314 \text{J/(mol} \cdot \text{K)}$, calculate the speed of sound and the root-mean-square speed of the molecules, and show that their ratio equals $\sqrt{\frac{\gamma}{3}}$.

Solution

(a) The light from almost every distant galaxy arrives with its spectral lines shifted towards longer wavelengths, the red end of the spectrum. This redshift is a Doppler effect: the galaxy is moving away from the Earth, so the wavelengths of the light it emits are stretched before they reach us. Almost all distant galaxies show this shift because the universe itself is expanding, carrying every galaxy away from every other. From the direction of the shift the astronomer learns whether the galaxy is receding, shown by a shift to the red, or approaching, shown by a shift to the blue. From the size of the shift, expressed as the fractional change in wavelength, the astronomer calculates the speed of the galaxy along the line of sight, since that fractional change equals the ratio of the recession speed to the speed of light.

(b) Convert the molar mass: $M = 28.0 \text{g/mol} = 0.0280 \text{kg/mol}$.
Speed of sound:

$$v_{\text{sound}} = \sqrt{\frac{\gamma RT}{M}} = \sqrt{\frac{1.40 \times 8.314 \text{J/(mol} \cdot \text{K)} \times 300 \text{K}}{0.0280 \text{kg/mol}}} = 353 \text{m/s}$$

Root-mean-square speed of the molecules:

$$v_{\text{rms}} = \sqrt{\frac{3RT}{M}} = \sqrt{\frac{3 \times 8.314 \text{J/(mol} \cdot \text{K)} \times 300 \text{K}}{0.0280 \text{kg/mol}}} = 517 \text{m/s}$$

Their ratio:

$$\frac{v_{\text{sound}}}{v_{\text{rms}}} = \frac{353 \text{m/s}}{517 \text{m/s}} = 0.683$$

Compare with $\sqrt{\frac{\gamma}{3}} = \sqrt{\frac{1.40}{3}} = \sqrt{0.467} = 0.683$.

Hence the ratio of the two speeds equals $\sqrt{\frac{\gamma}{3}}$, as required.

Example 42

(a) Explain why a person hearing a recording of his own voice for the first time finds that it does not sound like the voice he hears when he speaks, even though the recording is faithful.

(b) A car travels at 20m/s directly towards a tall vertical cliff while sounding a horn of true frequency 400Hz. Taking the speed of sound in air as 340m/s, calculate the beat frequency the driver hears.

Solution

(a) The quality, or timbre, of a voice is fixed by the particular mixture of harmonics it contains, not by its fundamental frequency alone. When a person speaks, the sound reaches his own ears by two routes. One route is through the air to the outer ear; the other is directly through the bones of his skull, which conduct the lower harmonics especially well and so make the voice in his head sound deeper and fuller. A recording captures only the sound that travelled through the air, which carries a larger proportion of the higher harmonics. The recording is faithful to that airborne sound, but it lacks the bone-conducted low harmonics the speaker is used to hearing, so the harmonic mixture is different and the voice sounds unfamiliar to him, even though everyone else hears him exactly as the recording does.

(b) The driver hears two tones at once. The first is the sound coming directly from the horn; because the horn and the driver move together with the car, they are at rest relative to each other and the driver hears this direct sound at the true frequency, 400Hz.

The second tone is the echo from the cliff, shifted twice. First the cliff receives the sound from the approaching car:

$$f_1 = \frac{fv}{v - u}$$

The cliff then reflects f_1, and the driver approaches the cliff as an observer:

$$f_{echo} = \frac{f_1(v + u)}{v} = \frac{fv}{v - u} \times \frac{v + u}{v} = \frac{f(v + u)}{v - u}$$

$$f_{echo} = \frac{400\text{Hz} \times (340\text{m/s} + 20\text{m/s})}{340\text{m/s} - 20\text{m/s}} = 450\text{Hz}$$

The beat frequency is the difference between the echo and the direct sound:

$$f_{beat} = f_{echo} - f = 450\text{Hz} - 400\text{Hz} = 50\text{Hz}$$

The driver hears 50 beats per second.

Example 43

(a) Explain why an X-ray can be used to photograph a broken bone inside the body while ordinary visible light cannot, even though both are electromagnetic waves.

(b) A tuning fork sounded with a sonometer wire of vibrating length 25.0cm produces 3 beats per second. When the vibrating length is shortened to 24.0cm, with the tension unchanged, the wire still produces 3 beats per second. Calculate the frequency of the tuning fork.

Solution

(a) Whether an electromagnetic wave passes through a material or is stopped by it depends on the wave and on the material. Visible light is strongly absorbed and reflected by the skin and soft tissue of the body, so it cannot reach a bone inside and carries back no information about it. X-rays have a much shorter wavelength and carry far more energy in each oscillation; they pass readily through soft tissue, which absorbs them only weakly, but they are absorbed much more strongly by the denser, heavier material of bone. An X-ray beam sent through a limb therefore emerges strongly weakened where it has crossed bone and barely weakened where it has crossed only flesh, and this difference, recorded on a detector, forms an image of the bone. Visible light cannot form such an image because it never reaches the bone in the first place.

(b) The fundamental frequency of a stretched wire at fixed tension is inversely proportional to its vibrating length, $f \propto \frac{1}{L}$. Shortening the wire from 25.0cm to 24.0cm therefore raises its frequency. The beat rate is the same magnitude before and after, so the wire must have been below the fork at 25.0cm and above the fork at 24.0cm. Writing f_0 for the fork frequency:

At 25.0cm the wire frequency is $f_0 - 3$

At 24.0cm the wire frequency is $f_0 + 3$

Since the wire frequency is inversely proportional to length:

$$\frac{f_0 - 3}{f_0 + 3} = \frac{L_2}{L_1} = \frac{24.0\text{cm}}{25.0\text{cm}}$$

$$25(f_0 - 3) = 24(f_0 + 3)$$

$$f_0 = 147\text{Hz}$$

The frequency of the tuning fork is 147Hz.

Example 44

(a) A daladala sounds its horn while travelling along the Bibi Titi road. Explain why a pedestrian standing on the verge hears the pitch of the horn change as the daladala passes, while a passenger seated inside the daladala hears no change of pitch at all, however fast the daladala travels.

(b) A steel railway rail on the TAZARA line is 1700m long. A single blow is struck on the rail at one end. Taking the speed of sound in steel as 5000m/s and in air as 340m/s, calculate the time interval between the two sounds heard by a worker with an ear pressed to the rail at the other end.

Solution

(a) The Doppler effect arises only when the source and the observer move differently through the air. The pedestrian on the verge is at rest in the air while the daladala carrying the horn moves; source and observer move differently, so the pedestrian hears the pitch rise as the horn approaches and fall as it recedes. The passenger seated inside the daladala, on the other hand, moves with the daladala, and so does the horn; the horn and the passenger share exactly the same motion through the air and are at rest relative to each other. With no difference in their motion through the medium there is no Doppler shift, and the passenger hears the horn at its true frequency however fast the daladala travels.

(b) The single blow sends a sound disturbance into the steel rail and also into the surrounding air. The sound through the steel covers the length of the rail at the speed of sound in steel:

$$t_{steel} = \frac{d}{v_{steel}} = \frac{1700\text{m}}{5000\text{m/s}} = 0.34\text{s}$$

The sound through the air covers the same distance at the speed of sound in air:

$$t_{air} = \frac{d}{v_{air}} = \frac{1700m}{340m/s} = 5.0s$$

The time interval between the two sounds is the difference:

$$\Delta t = t_{air} - t_{steel} = 5.0s - 0.34s = 4.66s$$

The worker hears the second sound, carried by the air, 4.66s after the first sound, carried by the rail.

These worked examples have carried the physics of sound, electromagnetic waves, and the Doppler effect through one situation after another, drawing the separate ideas of the chapter into a single connected picture. With that picture now clear, we are ready to put it to work. The Digging Deeper Exercise on the next page is where these ideas are tested, stretched, and made firmly our own.

DIGGING DEEPER EXERCISE 9

EXERCISE 9A: BINDER QUESTIONS

Question 1

A student states that an air particle in a sound wave is surrounded by the greatest pressure at the moment its displacement is greatest. Explain why this statement is incorrect.

Question 2

Explain why sound travels slightly faster through humid air than through dry air at the same temperature.

Question 3

Music played by a band reaches a distant listener in time and in tune, with no note arriving ahead of any other. Account for this observation.

Question 4

Explain why blowing harder across the mouth of a pipe eventually makes its note jump abruptly to a higher pitch, instead of rising smoothly.

Question 5

A student reads "0 dB" and concludes there is no sound at all. Account for what 0 dB represents and what a negative decibel level would mean.

Question 6

Two pure tones, of frequencies 50Hz and 3kHz, are produced at the same intensity. Explain why a listener does not hear them as equally loud.

Question 7

Explain why a sound wave cannot be polarised, whereas a light wave can.

Question 8

State and explain the phase relationship between the electric and magnetic fields of an electromagnetic wave travelling through a vacuum.

Question 9

Explain why the ultraviolet part of sunlight damages the skin, while the more intense visible light does not.

Question 10

In the Doppler effect for sound, explain why the motion of the source and the motion of the observer affect the observed frequency in different ways, even when their speeds are equal.

Question 11

A distant galaxy shows a large redshift. Explain whether this is better understood as motion of the galaxy through space, or as an expansion of the space between the galaxy and the Earth.

Question 12

A galaxy moves at high speed, but directly across the line of sight from the Earth. Explain why its spectral lines show no Doppler shift in wavelength.

EXERCISE 9B: REAL QUESTIONS

Question 13

Elephants in the Serengeti communicate with one another over distances of several kilometres using calls whose frequency lies below 20Hz. Explain why such low-frequency sound carries much further than ordinary speech.

Question 14

A wall between two rooms muffles the sound of a conversation next door, yet the deep bass of the music played there passes through clearly. Explain this difference.

Question 15

Kipanga: *"I can hear Mr. Akilikubwa clearly on the far side of that building, yet I cannot see him at all."*

Kipute: *"The same building stands in the way of both his voice and his image."*

Explain why the sound reaches Kipanga but the light does not.

Question 16

An echo from a cliff returns noticeably later on a cold morning than on a hot afternoon. Account for this difference.

Question 17

As a bucket is filled with water from a tap, the sound it makes rises steadily in pitch. Account for this rise in pitch.

Question 18

The whine of a mosquito is clearly heard in a quiet bedroom at night, but the same mosquito cannot be heard in a busy Kariakoo market by day. Account for this difference.

Question 19

A tall barrier built along the Morogoro Road reduces the traffic noise reaching the houses behind it, but never makes them completely silent. Explain why complete silence is not achieved.

Question 20

A mobile phone often loses its signal the moment its user steps inside a metal lift. Explain why the signal is lost.

Question 21

An AM radio crackles and pops during a thunderstorm, even when it is tuned to a strong station. Account for the crackling.

Question 22

A thermal-imaging camera can produce a clear picture of a person standing in complete darkness. Explain how this is possible.

Question 23

A television remote control must be pointed at the set and will not work from another room, yet a radio receiver plays normally far from the transmitter and behind walls. Explain this difference.

Question 24

Sealed medical instruments at a hospital are sterilised by exposing them to gamma rays. Explain why gamma rays are suitable for this purpose.

Question 25

Of the whole electromagnetic spectrum, the human eye is sensitive only to the narrow band of visible light. Explain why it is this band that the eye evolved to detect.

Question 26

A Doppler monitor in a clinic detects the heartbeat of an unborn baby. Explain how the movement of the baby's heart is detected.

Question 27

Kipanga: *"Since the universe is expanding, every galaxy must be moving away from us and must show a redshift."*

Kipute: *"Then explain the Andromeda galaxy, whose light is blueshifted."*

Account for the blueshift of the Andromeda galaxy.

EXERCISE 9C: HOT QUESTIONS

Question 28

On a cold night in Iringa the air temperature is 10°C, and on a hot afternoon it is 40°C. Calculate the percentage change in the speed of sound in air between the two.

Question 29

Carbon dioxide has $\gamma = 1.30$ and molar mass 44g/mol; hydrogen has $\gamma = 1.41$ and molar mass 2.0g/mol. Calculate the ratio of the speed of sound in carbon dioxide to that in hydrogen at the same temperature.

Question 30

One machine produces a sound level of 74dB at a certain point. Treating the machines as incoherent sources of equal intensity, calculate how many identical machines, running together, would raise the level at that point to 86dB.

Question 31

A point source radiates sound uniformly in the open air, producing a sound level of 90dB at a distance of 6.0m. Taking the threshold of hearing as $I_0 = 1.0 \times 10^{-12}\text{W/m}^2$, calculate the acoustic power output of the source.

Question 32

At a certain point in a sound wave the displacement amplitude is halved, the frequency remaining unchanged. Calculate the resulting change in the sound level at that point.

Question 33

A pulse of light crosses a glass block 0.30m thick, of refractive index 1.50. Taking the speed of light in vacuum as $c = 3.0 \times 10^8 m/s$, calculate the extra time the pulse takes compared with crossing the same thickness of vacuum.

Question 34

A hydrogen spectral line of laboratory wavelength 6563Å is observed from a distant galaxy at 6647Å. Taking $c = 3.0 \times 10^8 m/s$ and the Hubble constant as 70km/s per Mpc, calculate the recession velocity of the galaxy and its distance from the Earth.

Question 35

Taking the Hubble constant as $H_0 = 70 km/s$ per Mpc, and $1 Mpc = 3.09 \times 10^{19} km$, estimate the Hubble time in years and state what it represents.

Question 36

A whistle mounted on a truck is an open pipe of length 0.20m sounding its fundamental note. The truck travels at 30m/s directly towards a stationary roadside listener. Taking the speed of sound in air as 340m/s and neglecting end correction, calculate the frequency the listener hears.

Question 37

A train approaches a stationary observer at 25m/s, sounding a whistle of frequency 480Hz, on a day when the air temperature is 35°C. Taking the speed of sound in air at 0°C as 331m/s, calculate the frequency the observer hears.

Question 38

The cone of a loudspeaker, and the air in contact with it, executes simple harmonic motion of amplitude $1.0 \times 10^{-5} m$ at a frequency of 400Hz. Take the density of air as $1.2 kg/m^3$, the speed of sound as 340m/s, and $I_0 = 1.0 \times 10^{-12} W/m^2$. Calculate (a) the maximum speed of the air particles, (b) the pressure amplitude of the sound wave, (c) its intensity, and (d) the sound level produced.

Question 39

A closed organ pipe of length 0.25m is sounded together with a sonometer wire of vibrating length 0.50m and linear mass density $1.2 \times 10^{-3} kg/m$, the two being in unison at their fundamentals. Taking the speed of sound in air as 340m/s and neglecting end correction, calculate the tension in the wire.

Question 40

An ambulance siren of true frequency 700Hz approaches a stationary observer at 30m/s. Taking the speed of sound as 340m/s and treating the siren as a point source, calculate (a) the frequency the observer hears, and (b) the increase in the sound level the observer measures as the ambulance moves from 80m away to 20m away.

Question 41

In a resonance-tube experiment a tuning fork of frequency 512Hz gives consecutive resonances at air-column lengths of 16.0cm and 49.2cm. Taking the speed of sound in air at 0°C as 331m/s, calculate the speed of sound in the laboratory air, and hence deduce the air temperature.

Question 42

A satellite in low orbit transmits a radio beacon at a steady frequency of 150MHz, and approaches a ground station at 7.6km/s along the line of sight. Taking $c = 3.0 \times 10^8 m/s$, calculate the change in the received frequency, and explain why this shift falls to zero when the satellite passes directly overhead.

EXAMINATION QUESTIONS FOR PART TWO

Question 1

(a) A transverse pulse travelling along a stretched string is inverted when it reflects from a rigidly fixed end, but returns the same way up when it reflects from a free end. Account for this difference.
(b) In a 3D cinema the two lenses of the viewing glasses are polarising filters with their transmission axes set at right angles to one another. Account for how this lets each eye receive a separate image.
(c) A radio station in Dar es Salaam broadcasts at a frequency of 90.0MHz. Taking the speed of electromagnetic waves as 3.00×10^8m/s, calculate the wavelength of the broadcast and the time taken for the signal to reach a receiver 75km away.

Question 2

(a) Light passing through a wide doorway shows no noticeable spreading, while light passing through a slit a fraction of a millimetre wide spreads markedly beyond the straight-through direction. Account for this.
(b) A wave rolling in from deep water towards a gently shelving beach near Bagamoyo slows down and grows steeper as the water becomes shallower. Account for this.
(c) At a workshop one machine alone produces a sound level of 82dB at a worker's position, and a second machine alone produces 78dB at the same position. Calculate the sound level at that position when both machines run together. Take the threshold of hearing as $I_0 = 1.0 \times 10^{-12}\text{W/m}^2$.

Question 3

(a) The loudness of sound is measured on the decibel scale, which is logarithmic rather than linear. Account for the use of a logarithmic scale.
(b) Waves rolling into the harbour at Tanga strike the vertical harbour wall and travel back out, crossing the incoming waves and producing a choppy criss-cross pattern on the water. Account for this.
(c) Monochromatic light of wavelength 600nm falls normally on a diffraction grating ruled with 400 lines per millimetre. Calculate the highest order of diffraction maximum that can be observed.

Question 4

(a) Two transverse pulses travelling in opposite directions along a stretched string meet, overlap, and then continue on their way, each emerging with its original shape, size and speed. Account for this.
(b) A traffic police officer points a radar gun at a car approaching along the road. Account for how the gun is able to determine the car's speed.
(c) Plane-polarised light is incident on an analysing polaroid. Find the angle between the transmission axis of the polaroid and the plane of polarisation of the light for which the transmitted intensity is one quarter of the incident intensity.

Question 5

(a) An astronaut standing on the airless surface of the Moon would not hear a rock fall and shatter a few metres away. Account for this.
(b) The feathers of a peacock and the wing-cases of certain beetles show brilliant, shifting colours, yet contain no coloured pigment. Account for this.
(c) A transverse wave of amplitude 8.0mm and frequency 50Hz travels along a stretched string of linear mass density 0.020kg/m held under a tension of 80N. Calculate the speed of the wave and the average power it carries along the string.

Question 6

(a) A guitarist can change the frequency of the wave he sends along a string, but not the speed at which it travels. Account for the fact that the wave speed is fixed by the string while the frequency is fixed by the source.
(b) Looking at a distant streetlamp at night through the fine wire mesh of a window screen, an observer sees the single lamp spread into a regular pattern of bright spots. Account for this.
(c) An organ pipe closed at one end has a length of 0.55m. Taking the speed of sound in air as 340m/s, calculate the three lowest frequencies at which the pipe resonates.

Question 7

(a) A diffraction grating ruled with many thousands of lines produces much sharper and narrower bright lines than a double slit illuminated with the same light. Explain this.
(b) A bat hunting at dusk emits ultrasonic pulses and listens for the echoes. Explain how the bat can tell whether a moth ahead of it is flying towards it or away from it.
(c) Two waves of the same frequency travel in the same direction along a stretched rope. Their amplitudes are 3.0cm and 5.0cm and there is a constant phase difference of 60° between them. Calculate the amplitude of the resultant wave.

Question 8

(a) In a stationary wave on a stretched string, all the particles in one loop reach their extreme displacements at the same instant, while a particle in the next loop reaches its extreme on the opposite side of the string at that same instant. Explain.
(b) A military jet flying low and very fast over Dar es Salaam is followed by a sudden, loud bang heard at the ground. Account for this bang.
(c) Monochromatic light of wavelength 600nm passes through a single slit of width 0.20mm, and the diffraction pattern is observed on a screen 2.0m away. Calculate the angular position of the first dark fringe and the width of the central bright band on the screen.

Question 9

(a) In a two-slit interference pattern the bright fringes are all of very nearly the same brightness, but in a single-slit diffraction pattern the bright bands grow steadily fainter on either side of the centre. Account for this difference.
(b) Two children make a telephone from two empty tins joined by a length of string. When the string is pulled taut a child speaking into one tin is heard clearly at the other; when the string goes slack, nothing is heard. Account for this.
(c) A car travelling at 30m/s towards a vertical cliff sounds a horn of frequency 480Hz. Taking the speed of sound in air as 340m/s, calculate the frequency of the echo heard by the driver.

Question 10

(a) During a storm the flash of lightning is seen well before the thunder is heard, although both are produced at the same instant by the same lightning stroke. Account for the delay.
(b) A blacksmith heats a steel tool; as it cools in the air its polished surface passes through a sequence of colours, from pale straw through brown and purple to blue. Account for these colours.
(c) A wire fixed at both ends sounds a particular fundamental note, its length and mass per unit length kept constant. Determine the factor by which the tension must be increased to raise the frequency of the fundamental in the ratio 3:2.

Question 11

(a) A mass oscillating on a spring and a swinging pendulum bob appear to be very different systems, yet both execute simple harmonic motion. Account for the feature they share that makes this so.
(b) When a glass screen-protector is pressed onto a phone screen, patches of coloured rings and bands appear in the regions where the two surfaces almost, but not quite, touch. Account for these patterns.
(c) A motorcyclist rides at 20m/s directly towards a stationary siren that emits a steady note of frequency 660Hz. Taking the speed of sound in air as 340m/s, calculate the frequency of the note heard by the motorcyclist.

Question 12

(a) In a Newton's-rings experiment the dark and bright fringes are seen as a series of concentric circles centred on the point of contact. Account for the circular shape of the fringes.

(b) A diving board continues to move up and down for a short while after the diver has sprung off it. Account for this continued motion.

(c) A machine treated as a point source produces a sound level of 88dB at a distance of 2.0m in the open air. Calculate the sound level at a distance of 8.0m from the machine.

Question 13

(a) An organ pipe closed at one end sounds a note an octave lower than an open pipe of the same length. Account for this.
(b) A baby placed in a bouncer, a seat hung from a doorframe by an elastic strap, bounces up and down with a steady, regular rhythm. Account for this rhythmic bouncing.
(c) In a Young's double-slit experiment the bright fringes on the screen are 1.8mm apart when the apparatus is in air. The whole apparatus is then immersed in water of refractive index 1.33. Calculate the new spacing of the bright fringes.

Question 14

(a) If the amplitude of a simple harmonic oscillator is doubled, its maximum speed and its maximum acceleration both double, yet its period is unchanged. Account for this.
(b) In the great domed hall of a certain building, a person whispering quietly at one particular point can be heard clearly by someone far away at another particular point, though listeners between the two hear nothing. Account for this.
(c) A soap film of refractive index 1.33 and thickness 1.10×10^{-7}m is illuminated normally by white light. Calculate the wavelength in the visible range that is most strongly reflected.

Question 15

(a) Sound travels faster through warm air than through cold air. Account for this.
(b) Seen through a windscreen covered with fine scratches left by the wipers, the headlights of oncoming cars at night appear drawn out into bright streaks and spikes of light. Account for this.

(c) A 0.60kg mass hung from a single light spring oscillates vertically with a period of 0.80s. The mass is then hung from two such springs side by side, both supporting it together. Calculate the period of the vertical oscillations of the mass on the pair of springs.

Question 16

(a) Account for the fact that the total energy of a simple harmonic oscillator is proportional to the square of its amplitude.
(b) When a thin straight wire is held a short distance in front of one eye and a distant bright vertical lamp is viewed past its edge, a series of fine light and dark bands is seen. Account for these bands.
(c) An organ pipe open at both ends is to sound a fundamental note of frequency 256Hz. Taking the speed of sound in air as 340m/s, calculate the length of the pipe and the frequency of its second harmonic.

Question 17

(a) In Young's double-slit experiment a single narrow slit is placed between the lamp and the double slit. Account for the need for this single slit.
(b) A person who hears a recording of their own voice for the first time usually finds that it sounds unfamiliar and somewhat thinner than the voice they are used to hearing themselves speak in. Account for this.
(c) A 0.30kg mass hangs from a vertical spring of force constant 30N/m and is set oscillating with an amplitude of 0.080m. Calculate the period of the oscillation and the maximum and minimum tension in the spring during the motion. Take $g = 9.8\text{m/s}^2$.

Question 18

(a) Account for the fact that the kinetic energy of a simple harmonic oscillator varies with time at twice the frequency of its displacement.
(b) A doctor's stethoscope makes the faint sounds of a patient's heartbeat loud enough to hear clearly. Account for how it does this.
(c) Monochromatic light falls normally on a diffraction grating ruled with 500 lines per millimetre. The second-order maximum is observed at an angle of 36.0° to the straight-through direction. Calculate the wavelength of the light.

Question 19

(a) Sunlight reflected from a smooth water surface is found to be partially polarised, and at one particular angle of incidence it is completely polarised. Account for these observations.
(b) A heavy load hangs from the long cable of a crane. When the load is set swinging, the operator notices that paying out more cable makes the swinging slower and lazier. Account for this.

(c) A radar station sends out a pulse of microwaves and receives the echo reflected from an aircraft 1.20×10^{-4}s later. Taking the speed of electromagnetic waves as 3.00×10^8m/s, calculate the distance of the aircraft from the station.

Question 20

(a) Radio waves, visible light and X-rays are put to very different uses and have very different effects, yet they are all regarded as the same kind of wave. Account for this.
(b) Looking through polarised sunglasses at the toughened glass rear window of a car in daylight, a passenger sees a pattern of dark spots and blotches spread across the glass. Account for this.
(c) Calculate the length of a simple pendulum that beats seconds, that is, has a period of 2.0s, at a place where $g = 9.8\text{m/s}^2$. If this pendulum is then lengthened by 1.0%, calculate the resulting percentage change in its period.

Question 21

(a) Explain what is meant by the natural frequency of an oscillating system, and state the physical quantities that fix it for a mass on a spring and for a simple pendulum.
(b) At a football match the spectators rise and sit in turn, sending a 'wave' that travels right round the stadium, although no spectator leaves their seat. Account for the fact that the wave travels round the ground while the spectators do not.
(c) A girl standing in front of a high cliff claps her hands once and hears the echo 1.5s later. Taking the speed of sound in air as 340m/s, calculate her distance from the cliff. She then walks 60m directly towards the cliff and claps again; calculate the time after which she now hears the echo.

Question 22

(a) Explain why sound travels faster through a solid such as steel, and through a liquid such as water, than it does through air.
(b) A bungee jumper, after the first plunge, bounces up and down several times, each bounce smaller than the one before, until at last she hangs almost at rest. Explain this behaviour in terms of oscillations.
(c) Water waves travelling across a pond move at 1.8m/s and pass a fixed post 2.5 times every second. Calculate the wavelength of the waves, and determine the phase difference between two points lying 0.30m apart along the direction in which the waves travel.

Question 23

(a) A wave travelling along a light rope reaches a join where the light rope is tied to a much heavier rope. Explain why part of the wave is reflected at the join and part is transmitted into the heavier rope.
(b) A person standing near the corner of a building can hear someone speaking just around the corner, even though the speaker cannot be seen. Explain how the sound reaches the listener.
(c) A particle executing simple harmonic motion has a maximum speed of 0.30m/s and a maximum acceleration of 4.5m/s^2. Calculate the angular frequency, the period and the amplitude of the motion.

Question 24

(a) During simple harmonic motion the displacement, the velocity and the acceleration of the body all vary with time. Explain how the velocity and the acceleration are each related, in phase, to the displacement.
(b) A survey ship finds the depth of water beneath it by sending a pulse of sound down to the sea-bed and detecting the pulse that returns. Account for the way this method reveals the depth.
(c) A long string has a mass per unit length of 2.0×10^{-3}kg/m and is held under a tension of 50N. A vibrator at one end sends transverse waves of frequency 100Hz along it. Calculate the speed of the waves on the string and their wavelength.

Question 25

(a) The note of the siren of an ambulance is heard to fall in pitch at the moment the ambulance passes a stationary listener. Explain why the pitch falls.
(b) When a stone is dropped into a still pond, circular ripples spread outwards and grow steadily weaker the further they travel from the point where the stone entered. Account for the weakening of the ripples as they spread.
(c) A 0.25kg mass hung from a light spring oscillates vertically with a period of 0.50s. Calculate the force constant of the spring, and determine the further mass that must be added to the 0.25kg so that the period becomes 1.0s.

Question 26

(a) Use Huygens' principle to explain how a wavefront advances through a medium.
(b) A musician's metronome marks time with a steady tick, and the musician changes the tempo by sliding the small weight up or down its swinging arm. Explain, in terms of oscillations, why moving the weight changes the rate of ticking.
(c) The speed of sound in air is 340m/s when the air temperature is 15°C, that is 288K. Calculate the speed of sound when the air warms to 35°C. Hence calculate the wavelength in air of the note of a 480Hz tuning fork at each of these two temperatures.

Question 27

(a) Distinguish between the loudness and the pitch of a musical note, and state the physical property of the sound wave on which each of them depends.
(b) A heavy adult and a light child sit in turn on the same playground swing. Explain why each of them completes one full to-and-fro swing in the same time.
(c) Transverse wave pulses travel at 40m/s along a long cable stretched between two fixed posts 100m apart. A pulse is sent from one post; calculate the time after which the pulse, reflected from the far post, returns to its starting point. The tension in the cable is then increased to four times its original value; calculate the new return time.

Question 28

(a) Explain why the wavefronts produced by a small source of waves are very nearly circular close to the source, but appear practically straight to an observer far away from it.
(b) In a hospital, an image of an unborn baby is formed using ultrasound rather than ordinary audible sound. Suggest why ultrasound is used for this purpose.
(c) A body of mass 0.50kg executes simple harmonic motion with a total energy of 0.16J and an amplitude of 0.080m. Calculate the force constant of the system and the maximum speed of the body.

Question 29

(a) Explain why the amplitude of a forced oscillation becomes very large when the frequency of the driving force is close to the natural frequency of the system.
(b) A voice in an empty room with bare walls and floor rings on and sounds hollow, but once the room is furnished with curtains, carpets and soft chairs the hollow ringing disappears. Account for this difference.
(c) A stretched string vibrating as a stationary wave has a distance of 0.12m between a node and the antinode next to it, and the string vibrates 25 times each second. Calculate the wavelength of the waves on the string and the speed at which they travel along it.

Question 30

(a) Explain why an echo can be heard as a sound separate from the original only when the reflecting surface is sufficiently far from the listener.

(b) When two children turn a skipping rope steadily, the rope settles into a single broad loop, with the ends almost still and the middle sweeping out the widest path. Explain why this pattern forms.

(c) A particle moving with simple harmonic motion has a speed of 8.0cm/s when it is 6.0cm from the centre of its path, and a speed of 6.0cm/s when it is 8.0cm from the centre. Calculate the amplitude and the period of the motion.

Question 31

(a) Explain why a body executing simple harmonic motion spends, over each complete cycle, more of its time near the extremes of its path than near the centre.

(b) Boats moored in the sheltered water behind a harbour breakwater are still seen to rock gently, although the breakwater stands directly between them and the open sea. Account for the gentle rocking of the boats.

(c) Monochromatic light of wavelength 590nm is incident normally on a diffraction grating, and the first-order bright image is seen at an angle of 18.0° to the straight-through direction. Calculate the spacing of the grating lines and the number of lines per millimetre of the grating.

Question 32

(a) In a two-slit interference pattern there are dark fringes at which light arrives from both slits and yet no light is seen. Explain how light from the two slits can combine to give darkness, and account for what becomes of the light energy at a dark fringe.

(b) A hacksaw blade clamped firmly to the edge of a bench and twanged vibrates with a steady buzz; sliding more of the blade out beyond the edge of the bench makes the buzz deepen. Explain this behaviour in terms of oscillations.

(c) A wire 0.90m long is fixed at both ends and is made to vibrate in its third harmonic, sounding a note of frequency 200Hz. Calculate the wavelength of the waves on the wire and the speed at which they travel along it.

Question 33

(a) Explain why a stationary wave has points, the nodes, at which the medium never moves at all, even though each of the two waves that combine to form it would, on its own, move every point of the medium.

(b) A tuning fork sounds the very same musical note whether it is struck hard or struck gently. Explain why striking it harder changes the loudness of the note but not the note itself.

(c) In a Young's double-slit experiment the two slits are 0.40mm apart and the fringes are observed on a screen 1.5m away. The bright fringes are found to be 2.1mm apart. Calculate the wavelength of the light used.

Question 34

(a) A diffraction grating and a glass prism both spread white light into a spectrum, but the grating bends red light most and violet least, while the prism does the reverse. Account for the fact that the grating spreads the colours in the opposite order to the prism.

(b) A child pushes and pulls the end of a long springy coil lying on the floor, sending a pulse of compression running along its length to the far end. Account for the way the pulse travels along the coil.

(c) A spring of force constant 18N/m is cut into two equal pieces. Calculate the force constant of each piece, and the period of vertical oscillation of a 0.40kg mass hung from one of the pieces.

Question 35

(a) Explain why a mass suspended from a vertical spring oscillates with exactly the same period as it would if the same spring and mass were laid horizontally, even though gravity pulls on the hanging mass.

(b) On a misty night a bright street lamp is seen surrounded by a faint set of coloured rings. Suggest why these coloured rings appear.

(c) A continuous train of waves of frequency 20Hz travels along a stretched cord at a speed of 8.0m/s. Calculate the wavelength of the waves, and the number of complete waves contained in a 12m length of the cord.

Question 36

(a) Explain why the energy carried past a point each second by a progressive wave depends on the square of the amplitude of the wave.

(b) When a car approaches from far away at night, its two headlamps appear at first as a single point of light, and only when the car is much nearer do they become distinguishable as two separate lamps. Suggest why the two lamps cannot be told apart until the car is close.

(c) A particle executes simple harmonic motion with an amplitude of 0.10m. Calculate the displacement at which its kinetic energy is three times its potential energy, and the fraction of the total energy that is kinetic when the displacement of the particle is 0.080m.

Question 37

(a) A mass oscillating on a spring is replaced by a mass twice as large on the same spring. Explain why the period of oscillation increases by a factor of $\sqrt{2}$ rather than doubling.

(b) When a whip is cracked, the wave sent travelling down it moves faster, and with greater amplitude, as it nears the thin tip. Account for the speeding-up of the wave as it travels towards the tip.

(c) Unpolarised light of intensity 16W/m^2 is incident on a pair of polaroids whose transmission axes are set at 60° to each other. Calculate the intensity of the light transmitted through the first polaroid, and the intensity of the light emerging from the second.

Question 38

(a) Account for the equal spacing of the bright fringes across the screen in a Young's double-slit interference pattern.

(b) A child bounces steadily up and down on a pogo stick, a vertical toy with a strong spring for hopping, each bounce taking the same time as the one before. Explain, in terms of oscillations, what makes the bouncing so regular.

(c) A transverse wave of amplitude 1.5cm and frequency 12Hz travels along a stretched string at a speed of 9.0m/s. Calculate the maximum speed of a particle of the string, and state how it compares with the speed of the wave itself.

Question 39

(a) Explain why a train of water waves changes its direction of travel when it crosses, at an angle, into a region where the water is shallower and the waves move more slowly.

(b) A small boat at sea, rocked to one side by a passing wave, continues to roll to and fro with a steady rhythm of its own after the wave has gone. Account for this rolling.

(c) In a Newton's-rings experiment the plano-convex lens has a radius of curvature of 1.2m. The radius of the eighth dark ring, counted from the centre, is measured to be 2.4mm. Calculate the wavelength of the light used.

Question 40

(a) Explain why the sky appears blue to an observer on the ground during the day.

(b) Waves approaching a straight beach are often seen to arrive very nearly parallel to the shoreline, even when, far out, they were travelling towards the shore at a noticeable angle. Explain why the waves arrive nearly parallel to the shore.

(c) A particle moving with simple harmonic motion has a period of 0.40s and travels a total distance of 24cm in each complete oscillation. Calculate the amplitude of the motion, its maximum speed and its maximum acceleration.

PART THREE

ELECTRICITY

Having traced the behaviour of heat in Part One and the motions and waves that carry energy through matter and space in Part Two, we now turn to one of the most far-reaching forces in nature: the electric force. Electricity examines charges at rest and in motion, beginning with electrostatic fields and potential, moving through the storage of charge in capacitors, the steady flow of current in conductors and circuits, and finally the alternating currents that power modern life. These ideas explain how charge gives rise to forces and fields, how energy is stored and delivered through electrical networks, and how the same principles govern systems ranging from a spark across a doorknob to national power transmission, radio tuning, and electronic devices. Mastering this part will deepen your understanding of charge, field, and circuit behaviour, and prepare you for more advanced studies in electronics, electromagnetism, and applied technology.

Chapter 10

ELECTROSTATICS

INTRODUCTION

Mid-afternoon at Miono Secondary School. The corridor outside the physics laboratory has been polished so aggressively that students walk through it like people crossing a frozen lake. For two weeks the brass door handle has been famous. The Form Sixes call it *the snake* because every afternoon somebody touches it and jumps backwards with wounded dignity.

Today it is Kipanga's turn.

He drags his black rubber shoes slowly across the polished floor while several Form Sixes watch from a safe distance with the excitement normally reserved for football penalties and public disasters. Kipanga approaches the brass handle carefully and stretches out one finger.

Nothing happens.

The Form Sixes groan in disappointment.

Kipanga frowns, walks back to the far wall, and rubs his shoes across the corridor again, harder this time. Even the bougainvillea beside the laboratory seems interested. He returns to the handle and lowers his finger toward it with the caution of a man touching a sleeping cobra.

SNAP!

A tiny blue spark jumps from the brass before he even touches it.

Kipanga yelps and nearly throws himself into the notice board while the Form Sixes explode with laughter.

"The snake has bitten another one!"

Mwajuma, passing with a flask from the chemistry prep room, jumps so badly she almost drops it. Kipute, wearing open sandals and carrying absolutely no fear, walks straight to the same handle, grabs it, opens the door, and walks inside unharmed.

The Form Sixes look cheated.

From the louvre window, Mr. Akilikubwa finally steps into the corridor with the calm expression of a man who has been waiting all week for this exact foolishness.

Mr. Akilikubwa: *"Kipanga, congratulations. You have just produced several thousand volts using only your shoes and poor decision-making."*

Kipanga (staring at his finger suspiciously)**:** *"Sir, there is no socket anywhere near me."*

Mr. Akilikubwa: *"You do not need a socket. Every step moved electrons from the floor onto your body. Your rubber shoes trapped the charge instead of letting it escape. By the time you reached the brass handle, your finger and the metal had a huge potential difference between them. The air finally gave up and the charge jumped across."*

Kipute (lifting one sandal)**:** *"Then why did the snake ignore me?"*

Mr. Akilikubwa: *"Because your sandals leak charge away continuously. Kipanga stored charge like a battery. You leaked it away like a cracked bucket."*

Mwajuma (edging closer)**:** *"Sir… is that the same physics as lightning?"*

Mr. Akilikubwa (pointing upward dramatically)**:** *"Exactly the same physics. A cloud stores charge the way Kipanga did. When the electric field becomes strong enough, the air breaks down and the charge jumps through the sky. Lightning is simply the giant version of Kipanga embarrassing himself in public."*

The Form Sixes cheer loudly at this scientific interpretation.

Kipanga rubs his finger carefully.

Kipanga: *"Sir, so the snake is just electrostatics?"*

Mr. Akilikubwa (smiling)**:** *"Correct. By the end of this chapter, you will calculate the voltage in that spark, the electric field in the air gap, and the energy released into Kipanga's unfortunate finger. You will also discover that electrostatics powers photocopiers, removes smoke from factory chimneys, and occasionally attacks Form Five students in corridors."*

Kipute is already opening her notebook.

Kipanga, meanwhile, keeps a respectful distance from the brass handle.

Welcome to Chapter 10. We are studying charges at rest: where they come from, what forces they exert on one another, what invisible field they fill the space around them with, what potential and potential energy they carry, what happens when they meet a perfect conductor, and what they do when an outside agent pushes them around.

Two debts from earlier chapters are paid here. The first is to ***Foundations of Physics***, Chapter 7. Newton's law of universal gravitation,

$$F = \frac{Gm_1m_2}{r^2}$$

had a twin all along, and you are about to meet it as Coulomb's law,

$$F = \frac{kQ_1Q_2}{r^2}$$

Everything that worked for planets carries across with the labels changed.

The second debt is to Chapter 1 of this book. The U(r) curve you met for two molecules will reappear as the U(r) of two point charges, which is the cleanest version of that picture you will ever see, with one smooth curve and a sign that comes from the charges.

We close the chapter the way we opened it, with a static-electric story. By the last pages the brass handle, the photocopier, the lightning conductor, the smoke at Tanga Cement, and the cassava-yard conveyor will all sit on one page of running prose, each explained by the same handful of equations. By that page Kipanga will know exactly which electrons crossed the air gap this afternoon, why, and at what speed.

Let us begin.

ELECTRIC CHARGE

Two thousand six hundred years ago, a Greek philosopher called Thales noticed that a piece of amber rubbed with fur would pick up small bits of dry grass and chaff. The Greek word for amber is *elektron*, and that single observation has, by a chain of inheritances, given us the word electricity and the entire industry that lights our towns, charges our phones, and runs the door-handle spark that opened this chapter. Thales did not know what was happening; nobody did, for the next two thousand four hundred years. But the experiment is real, repeats every time, and is the door into a body of physics that this chapter is about to open.

Before we can write down a single equation about electric forces, we have to understand what the rubbed amber actually has more of, and the rubbed fur actually has less of, after the rubbing. That common quantity, the one that flows from one to the other, is what we call **electric charge**. This first section answers four questions about it. *What kinds of charge are there? Can charge be created or destroyed, or does it only move around? Is charge a smooth quantity like length, or does it come in fixed indivisible packets like a brick wall?* And finally, *by what physical processes does a neutral body become a charged one?* Once those four questions are answered, the machinery of Coulomb, of fields, of Gauss and of potential will follow without resistance.

The two kinds of charge

Take two glass rods, rub each of them with a piece of silk, and hang one of them from a thread so it can swing freely. Bring the second glass rod close to the first. The hanging rod swings away. The two rubbed glass rods repel each other. Now repeat the experiment, but with two plastic rods rubbed with fur instead of two glass rods rubbed with silk. The two rubbed plastic rods also repel each other. Finally, bring a rubbed plastic rod close to the rubbed glass rod hanging on its thread. The hanging rod swings towards the plastic rod, not away from it. Glass and plastic attract.

This humble experiment, which any reader can perform in a school laboratory in fifteen minutes, contains the first big result of electrostatics. The rubbing has put something onto each rod, and that something comes in exactly two flavours. Two glasses repel, two plastics repel, but glass and plastic attract. There is no third flavour and there is no neutral flavour, because all the rods involved have been rubbed. We name the two flavours after the signs we give them on paper. The charge that builds up on the glass rod rubbed with silk is called **positive charge.** The charge that builds up on the plastic rod rubbed with fur is called **negative charge.** The choice of which one to call positive was made by Benjamin Franklin in the eighteenth century and is, in honesty, an accident of history; if Franklin had flipped his coin the other way, the electrons in your phone would today be the positive carriers. The physics would be unchanged.

The two-flavour rule that the experiment uncovered can now be stated as a working principle of the chapter. *Like charges repel each other, and unlike charges attract each other.* That sentence will be the seed from which every Coulomb force in the chapter grows, and the rubbed glass rod hanging on its thread is the diagnostic test that decides, in any new experiment, whether the unknown sample is positively charged, negatively charged, or not charged at all. If the unknown attracts the glass, the unknown is negative. If the unknown repels the glass, the unknown is positive. If the unknown does nothing measurable to the glass, the unknown is neutral.

Conservation of charge

Where did the charge on the glass rod come from? The glass rod was neutral before the rubbing, the silk was neutral before the rubbing, and yet after a few seconds of contact the glass had become positive. Either charge was made out of nothing during the rubbing, which would be an extraordinary thing to claim, or some charge that was already there moved from one

object to the other. Careful experiments performed over the last three centuries have settled the question decisively in favour of the second answer. The silk did not create the charge on the glass. The silk supplied it.

In every experiment that has ever been performed, the total electric charge of an isolated system, that is, of a system with no exchange of matter with anything outside it, has been observed to remain constant. Charge can move from one body to another, can be hidden inside a neutral object by being equal and opposite to an unseen partner, and can be moved around at great speed inside wires; but it cannot be created and it cannot be destroyed. The total stays the same forever.

This experimental fact has been elevated, by its perfect record, to the status of a named principle. We call it the **principle of conservation of charge:**

In any isolated system the algebraic sum of the positive and negative charges is constant, no matter what physical or chemical processes take place inside the system.

When the glass rod was rubbed with silk, the silk gave up some of its electrons to the glass, and ended that exchange with a deficit of electrons of exactly the same magnitude as the deficit the glass had gained. The glass rod is positive by exactly as many micro-coulombs as the silk is negative; the books balance.

A second, more dramatic illustration is the reaction in which a single high-energy gamma photon, carrying no charge at all, materialises into an electron of charge minus e and a positron of charge plus e. The total charge before the event is zero. The total charge after the event is also zero, because the electron and positron carry equal and opposite charges. Charge has not been made out of nothing, even in this most exotic process; it has been made out of an equally large pair that adds to zero. The book of electric accounts is never, ever overdrawn.

Quantisation of charge

A neutral piece of glass can be rubbed against silk to give it a charge of two microcoulombs. *Can it then be rubbed a little less to give it a charge of, say, exactly one half of a microcoulomb? Or one third? Or any tiny fraction at all?* The answer, discovered in 1909 by the American physicist Robert Millikan in his famous oil-drop experiment and confirmed by every careful measurement made since, is **no.** There is a smallest unit of free charge in nature, and the charge on any isolated object is always a **whole-number** multiple of that smallest unit.

The smallest free charge is the magnitude of the charge on a single electron, and is given the symbol **e**. Its measured value, accurate to many decimal places, is $e \approx 1.6\times10^{-19}$C. An electron carries a charge of $-e$, that is -1.6×10^{-19}C, and a proton carries a charge of $+e$, that is $+1.6\times10^{-19}$C. Every free electric charge that has ever been measured is an integer number n of these elementary units, and the total charge Q on an object that has gained or lost n electrons is therefore:

$$\mathbf{Q = ne},$$

where n is a positive whole number when the body has lost electrons (and is therefore positive) and a negative whole number when the body has gained electrons (and is therefore negative). This property, that free charge always comes in whole-number multiples of e, is called the **quantisation of charge.** The word quantisation means that the quantity is not a smooth river but a stack of identical bricks, and you can have one brick, two bricks, a million bricks, but never one and a half bricks.

For ordinary laboratory objects the brick is so small that the staircase looks like a ramp. A charge of one microcoulomb is six and a quarter trillion electrons, so adding or removing one more electron is a step too small to notice. Quantisation matters very much, though, at atomic scales, where individual electrons are the actors. This is not as a peculiar rule that physicists invented but as the simple statement that you cannot have half an electron, just as you cannot have half a brick.

How a body becomes charged

A neutral body has equal amounts of positive and negative charge, the positives bolted onto the nuclei and unable to move, the negatives carried by the outer electrons of the atoms and able to move with relative ease through some materials. To charge a body is to give it a net excess of one kind, which by conservation requires that an equal amount of the other kind ends up somewhere else. Three standard procedures do this, namely: charging by friction, charging by conduction, and charging by induction. Each one is best understood by following the electrons.

1) Charging by friction

When two different insulating materials are rubbed firmly against each other, electrons from the surface atoms of one material are transferred mechanically to the surface atoms of the other. Which material gives up its electrons and which one receives them depends on the chemical nature of the two surfaces, and the order has been catalogued experimentally into a list called the **triboelectric series.** Near one end of the series sit *glass, hair, and wool,* which **lose electrons easily**. Near the other end sit *rubber, amber, and plastic*, which **accept electrons easily**. The **further apart** two materials sit on the series, *the more charge each one acquires when they are rubbed together.*

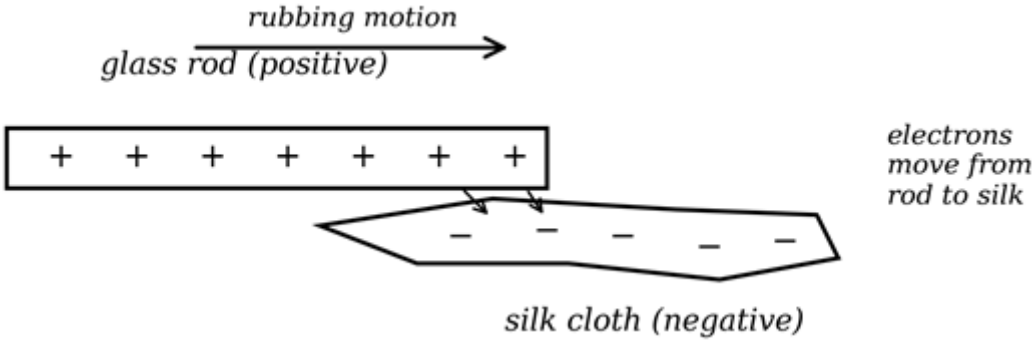


Figure: ***Charging by friction.*** *Rubbing a glass rod with silk transfers electrons from the surface of the glass to the surface of the silk. The glass rod is left with a deficit of electrons and is therefore positive; the silk is left with an excess of electrons and is therefore negative. The total charge of the rod-and-silk system is zero before and after the rubbing, as required by the conservation of charge.*

Charging by friction is the procedure by which the door handle of the introduction acquired its charge. Kipanga's rubber soles dragged across the polished concrete of the corridor and exchanged electrons with the floor; the shoes (and Kipanga's body, which is in good electrical contact with the shoes) became charged. The procedure works best on dry days. On a humid day the moisture in the air forms a thin conducting film over the surfaces, and the charge leaks away through that film as fast as the rubbing puts it on. This is why the famous door handle is famous in the dry season and unremarkable in the rainy one.

The reader who would like to see **triboelectric charging** in their own hand can do so in two minutes. Tear a sheet of writing paper into pieces no larger than a fingernail, scatter them on a dry tabletop, and pass a plastic comb a few times briskly through clean dry hair. Held two centimetres above the scattered paper, the comb will lift a handful of the pieces up to its teeth and hold them there. Hair sits above plastic on the triboelectric series, so the comb gains electrons during the rubbing and becomes negatively charged. The paper, although neutral overall, is polarised by the field of the comb (the same kind of induced separation we shall see in the next subsection), and the net force on each polarised piece is upward towards the comb. After a few seconds some of the lifted pieces fall back to the table, because the comb has lost a little of its charge to the surrounding air. On a humid day almost nothing rises; on a dry day the whole demonstration works first time.

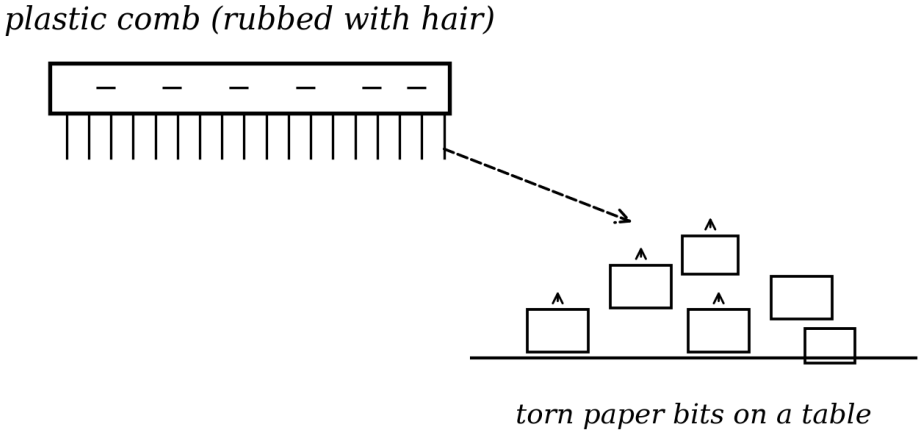


Figure: *A plastic comb that has been rubbed through dry hair gains a small negative charge and lifts torn pieces of paper from a flat surface. The paper itself has no net charge, but the comb induces a separation of charge inside each piece (negative side away, positive side towards the comb), and the resulting attraction is sufficient to lift the piece against gravity.*

2) Charging by conduction

A second way to charge a previously neutral body is to put it into direct contact with a body that is already charged. If the two bodies are conductors, the excess charge on the charged body redistributes itself between the two until both objects sit at the same electric potential, and they share the available charge between them. The previously neutral body now carries a charge of the same sign as the original charged body, by exactly the amount that has flowed across. The original body has lost an equal amount, and the total charge of the two-body system is unchanged. **Charging by conduction** therefore *always produces a body of the same sign as the body that did the charging.*

A small subtlety of conduction is worth flagging. *If the previously neutral body is large compared to the charged body that touches it, the charged body loses almost all of its charge to the larger one.* In the limit of a very large body, such as the Earth itself, the charged body loses essentially all of its charge and ends up neutral. This is why touching a charged metal object to

the ground discharges it almost completely, and it is the principle behind every earth wire in every electrical installation in your home.

3) Charging by induction

The third procedure is the cleverest, because it produces a charged body without the charging body ever touching it. It is called charging by induction, and it relies on the fact that the free electrons in a conductor are free to move when they feel an electric force.

The procedure has four steps and is best followed on the figure that accompanies this paragraph. In Step 1 a charged rod, say a positively charged glass rod, is brought close to a neutral conducting sphere mounted on an insulating stand. The sphere does not touch the rod. In Step 2 the free electrons inside the sphere, being negative, are attracted by the positive rod and crowd onto the near side of the sphere, leaving the far side with an equal deficit of electrons (that is, a positive surface region). The sphere as a whole is still neutral, but its charge has rearranged itself.

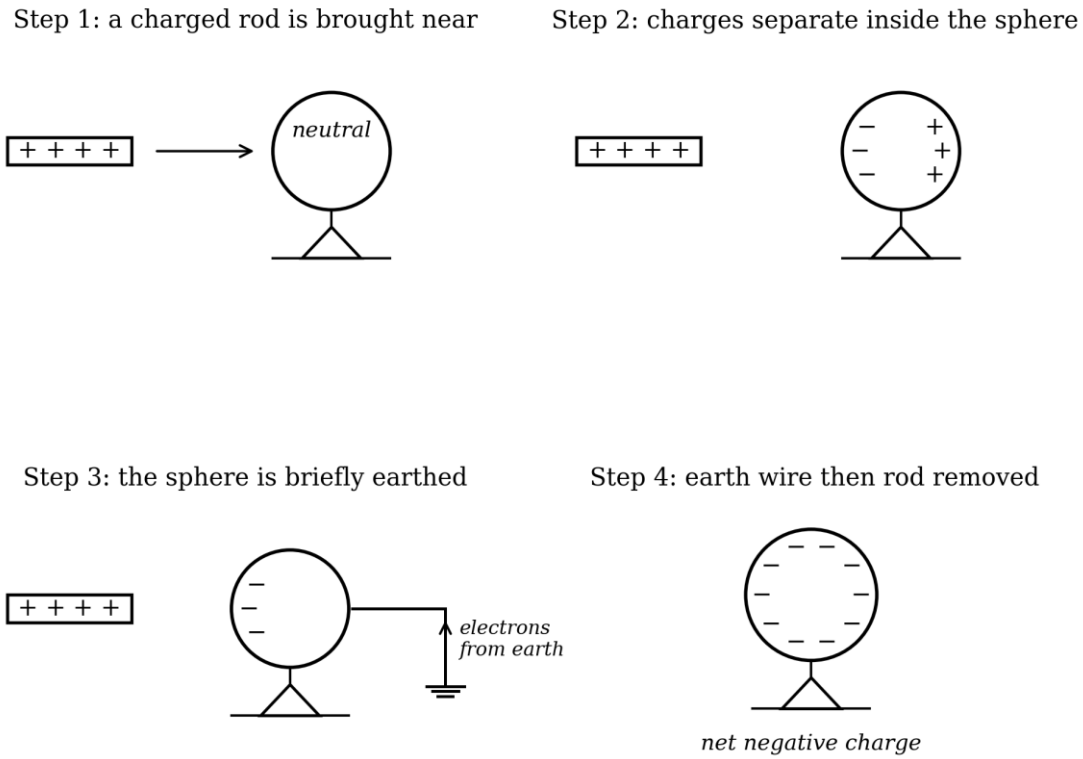


Figure: ***Charging by induction.*** *In Step 1 a positively charged rod is brought near a neutral conducting sphere on an insulating stand. In Step 2 the free electrons in the sphere are attracted to the near side, leaving the far side with an equal positive region. In Step 3 the sphere is briefly earthed, so the unwanted far-side positive region is neutralised by electrons that flow up from the earth. In Step 4 the earth connection is removed first and then the rod is removed, leaving the sphere with a net negative charge of the opposite sign to the rod.*

In Step 3 a connecting wire briefly joins the sphere to the earth. The earth is effectively a very large reservoir of charge, and electrons flow up the wire from the earth into the sphere to neutralise the positive far-side region. (One can equivalently say that the positive charges flow down the wire into the earth; the two pictures are equally valid and produce the same final state.) The sphere now carries a net excess of negative charge, all of it bound to the near side by the still-present positive rod. In Step 4 the earth wire is removed first, and only then is the rod removed. With the rod gone, the trapped electrons on the near side are free to redistribute, and they spread out over the entire surface of the sphere. The sphere is left with a net negative charge, opposite in sign to the rod that caused all the rearrangement.

Three small details of induction repay attention.

- The rod never touched the sphere, so no charge crossed between them; the rod's own charge is unchanged at the end of the procedure.
- The order of removal matters: the earth wire must be lifted before the rod, because lifting the rod first would simply let the bound electrons on the sphere escape back to earth before they can be trapped.
- The sphere ends up with a charge of the sign opposite to the rod, in contrast to charging by conduction which always produces the same sign.

Examination questions that ask students to predict the final sign always reward those who remember which procedure was used.

The book has so far been all words and no numbers, like a football commentator who never tells you the score. Two worked examples now arrive to fix that.

BINDER Example 1

A small metal sphere carries a positive charge of +5.0μC after being rubbed with silk. Calculate the number of electrons that the sphere has lost during the rubbing. Take the elementary charge to be $e = 1.6\times10^{-19}$C.

Solution

The charge on the sphere is positive, which by the quantisation rule means the sphere has lost a whole number of electrons rather than gained any. Calling that whole number n, the magnitude of the charge on the sphere is the product of n and the elementary charge:

$$Q = ne$$

Make n the subject and substitute:

$$n = \frac{Q}{e} = \frac{5.0 \times 10^{-6}\text{C}}{1.6 \times 10^{-19}\text{C}} = 3.125 \times 10^{13}$$

The sphere has therefore lost about 3.125×10^{13} electrons (about thirty-one trillion electrons).

Making Sense of the Answer: *Thirty-one trillion sounds like an unimaginable number to lose, and yet five microcoulombs is a perfectly ordinary amount of charge to find on a rubbed object in a school laboratory. The two facts together tell us that a single electron carries an extremely small charge, so an ordinary rubbed body has to lose a colossal number of electrons before the total charge becomes laboratory-sized. The sphere as a whole, after losing those thirty-one trillion electrons, has not lost any measurable mass and has not lost any measurable size.*

Think Like a Physicist: *Whenever you are given a charge in coulombs and asked for a count of electrons, the answer is always going to be the charge divided by e, and the only question is where the decimal point lands. The trick is to remember that e is a very small number, of order 10^{-19}, so dividing by it produces a very large count. Conversely, whenever you are asked for the charge produced by a given number of electrons, you will multiply by e, and the count will collapse into a tiny charge in coulombs.*

REAL Example 2

On a dry afternoon a child rubs an inflated rubber balloon briskly against her woollen jumper and presses it lightly against the painted plaster wall of her classroom. The balloon sticks to the wall and remains there for several minutes without any tape or glue. Explain in terms of charge conservation and induction:

(a) what sign of charge the balloon acquires from the rubbing,
(b) why the balloon sticks to a wall that is itself electrically neutral, and
(c) why the balloon eventually falls off after a long enough wait.

Solution

(a) Rubber sits well below wool on the triboelectric series, which means that when the two materials are rubbed together electrons are transferred from the wool to the rubber. The balloon, made of rubber, therefore gains electrons during the rubbing and ends up with a negative net charge. The jumper loses an exactly equal number of electrons during the same rubbing and ends up with a positive net charge of the same magnitude. The total charge of balloon plus jumper is zero before the rubbing and zero after, in obedience to the conservation of charge.

(b) The painted plaster wall is neutral overall, with equal amounts of positive and negative charge tied up in its atoms. When the negatively charged balloon is brought close to the wall, its negative surface charge pushes the negative electrons in the surface layer of the wall slightly inward, and pulls the positive nuclei of those atoms slightly outward. This is induction in an insulator: no charge is free to travel across the wall, but each atom is slightly distorted, so the side of every atom nearer the balloon becomes a little positive, and the side further from the balloon becomes a little negative. The result is a thin layer of induced positive charge on the wall at the spot directly under the balloon. This thin induced positive layer attracts the negative balloon. Because the induced positive layer is closer to the balloon than the induced negative layer behind it, the attraction is stronger than the repulsion, and the net force on the balloon is towards the wall. That net force is what holds the balloon up against its own weight.

(c) The induced separation of charge in the wall is not maintained by the wall itself; it is maintained by the continued presence of the charged balloon. Over time, two things happen. The first is that some of the negative charge on the balloon's surface slowly leaks away into the air, especially if the air carries a little humidity, and the balloon's charge falls. The second is that the polarised molecules in the wall slowly relax back to their unbiased positions, partly because of thermal motion and partly because their internal electric forces resist being held apart. As the balloon's charge falls and the wall's polarisation relaxes, the attractive force weakens until it is no longer large enough to support the balloon's weight, at which point gravity wins and the balloon slides down the wall and falls. On a humid day this happens in a shorter time (within a minute or so); on a dry day it takes longer time.

Making Sense of the Answer: *The balloon-on-the-wall trick, which most readers will have performed as small children long before they took a physics class, is in fact a beautifully complete electrostatics experiment. It uses friction to charge the balloon, induction to charge the wall, and the resulting attraction to hold the balloon up against gravity. Three of the four big ideas of this section (the two kinds of charge, conservation of charge, and induction) appear in a single household demonstration. The fourth, quantisation, is hidden inside the calculation of just how many electrons moved, and is the only one that requires a number to be revealed.*

Think Like a Physicist: *An attractive force between a charged object and an apparently neutral one is almost always a sign that induction is at work. The neutral object is responding to the charged one by separating its own internal charges slightly, and the separated charges then experience unequal forces because they are at slightly different distances from the charged object. You will meet this same idea three more times in the chapter, once for the rubbed comb earlier in this section, once for the bird on the high-voltage power line, and once for the electrostatic precipitator that cleans Twiga Cement's chimney exhaust.*

Two worked examples now lie behind us, and four big results have been bagged: that charge comes in two flavours, that the books always balance, that the smallest brick in the wall has the value of e, and that a body can be charged by three different procedures. The balloon you stuck to the wall in primary school turns out to have been a working electrostatics laboratory, and you can now charge for the rent.

With two charged objects in hand and a clear list of the rules that govern their charges, the next obvious question is the size of the force one of them exerts on the other. That question was answered, in elegant quantitative form, by a French military engineer in 1785. His result is the subject of the next section, and the inverse-square law it contains will turn out to be the electrostatic twin of the gravitational law you met in ***Foundations of Physics.***

COULOMB'S LAW AND THE SUPERPOSITION PRINCIPLE

The previous section settled what charge is, where it comes from, and how it moves between bodies. It did not say how much force two charged bodies exert on each other. The man who answered that question lived two hundred and forty years ago in revolutionary France, and the answer he found is one of those rare results in physics that fits on a single line, applies from the size of an atom to the size of a thunderstorm, and never has a known exception in any laboratory experiment. We meet his result, meet the law of superposition that turns it into a tool for arbitrarily many charges, and find, at the end of the section, that the law we have just learned has a twin in another part of this series of books.

What Coulomb measured

In 1785, the French military engineer Charles-Augustin de Coulomb hung a light horizontal insulating rod from a thin silk fibre inside a tall glass cylinder. At one end of the rod sat a small metal ball; at the other a small paper counterweight. Coulomb gave the ball a known electric charge, and then brought a second charged ball, fixed on its own insulating rod, very close to the first one without touching it. The repulsion between the two charged balls twisted the suspending fibre by a small angle, and Coulomb read the angle on a scale at the top of the cylinder. Because the twist of a thin fibre is known to be proportional to the torque applied to it, the read angle told him, directly, the size of the repulsive force.

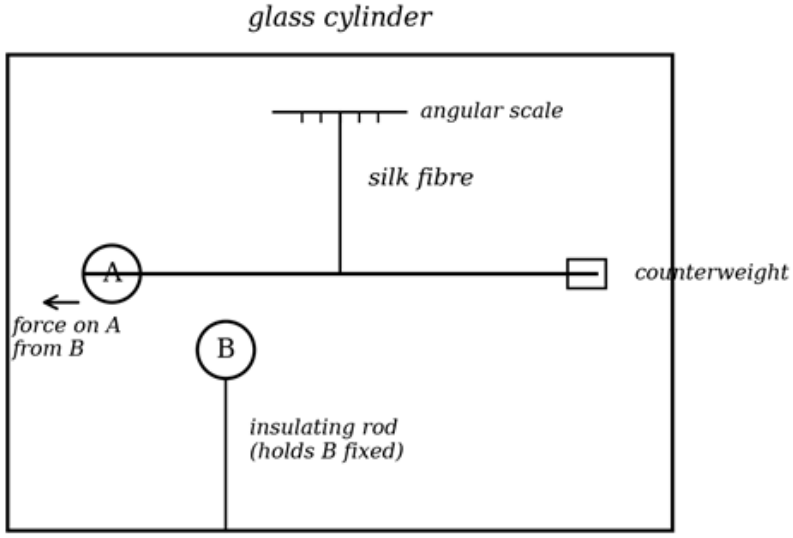


Figure: ***Coulomb's torsion balance.*** *A light insulating rod, suspended at its centre by a thin silk fibre inside a tall glass cylinder, carries a small charged ball A at one end and a paper counterweight at the other. A second charged ball B, fixed on its own insulating stem, is brought close to A. The repulsion between the two charged balls twists the fibre, and the angle of twist, read on the scale at the top, measures the force between them.*

Coulomb repeated the experiment many times. He doubled the distance between the balls and saw the force fall to one quarter of its earlier value. He tripled the distance and saw it fall to one ninth. The pattern was unmistakable. The force between two charged balls falls as the square of the distance between them increases:

$$F \propto \frac{1}{r^2}$$

Coulomb then varied the amount of charge on the two balls, using the trick of sharing charge by contact (a charged ball briefly touched to an identical neutral ball shares its charge equally between them, so dividing the original charge by two; repeating the contact divides by four; and so on). He found that the force is proportional to each of the two charges:

$$F \propto Q_1 Q_2$$

Both proportionalities are observed at the same time, so the force is jointly proportional to the product of the two charges and inversely proportional to the square of the distance between them. Combining the two relations into a single statement gives the inverse-square law of electrostatics, which is universally known today as **Coulomb's law:**

The electrostatic force between two stationary point charges is directly proportional to the product of their charges and inversely proportional to the square of the distance between them, acting along the line joining the charges.

Mathematically:

$$\mathbf{F = \frac{kQ_1Q_2}{r^2}}$$

The constant k is the proportionality constant whose numerical value depends on the units in which charge, distance, and force are measured. In the SI system, in which charge is measured in coulombs, distance in metres, and force in newtons, the experimentally measured value of this constant is:

$$k = 9.0 \times 10^9 \ Nm^2/C^2,$$

to two significant figures, and is the value we shall use throughout every numerical calculation in this chapter.

Why the constant is written as 1/(4πε₀)

In nearly every physics textbook that goes beyond an introductory level, the constant k in Coulomb's law is rewritten in the slightly odd-looking form:

$$k = \frac{1}{4\pi\varepsilon_0},$$

so that the law is written as

$$\mathbf{F = \frac{1}{4\pi\varepsilon_0}\frac{Q_1Q_2}{r^2}},$$

in which ε_0 is called the permittivity of free space. The first reaction on meeting this expression is to ask, very reasonably, why anyone would replace a single clean letter k with a fraction containing a 4π factor. The reason is geometric, and it foreshadows a powerful theorem we shall meet later in this chapter (Gauss's law).

Imagine a single point charge sitting at the centre of an imaginary sphere of radius r. The surface area of that sphere is $4\pi r^2$. If we picture the lines of electric influence that emerge from the charge as fanning out uniformly into all directions, then the density of those lines on the surface of the sphere is the total number of lines divided by $4\pi r^2$. Writing the constant in Coulomb's law as $\frac{1}{4\pi\varepsilon_0}$ folds this geometric 4π factor directly into the law, and leaves the new constant ε_0 to carry only the physics of how strongly empty space resists being threaded by electric influence. It is, in the end, an elegant book-keeping choice that pays its dividends in three sections' time, when Gauss's law uses exactly this rearrangement to make the algebra of fields and fluxes work in one line. For the moment, take it on trust, and remember the numerical value:

$$\varepsilon_0 = 8.854 \times 10^{-12} \ C^2/(Nm^2)$$

Permittivity and relative permittivity

The constant ε_0 that appears above is, strictly speaking, the permittivity of a vacuum: the space between the two charged balls of Coulomb's torsion balance contained almost no matter, and the value of ε_0 just quoted describes that empty space. If the two charges are immersed in a different material, the force between them changes, and almost always falls. A drop of water placed between two charged spheres reduces the force between them by a factor of about eighty, because the polar water molecules rearrange themselves to oppose the field, weakening it.

To capture this change quantitatively, we introduce a dimensionless number ε_r called the **relative permittivity** (also termed as the **dielectric constant**) of the material. The permittivity of the material is then:

$$\boldsymbol{\varepsilon = \varepsilon_r\varepsilon_0} \ \left(\text{from } \boldsymbol{\varepsilon_r = \frac{\varepsilon}{\varepsilon_0}}\right)$$

and Coulomb's law inside the material reads:

$$F = \frac{1}{4\pi\boldsymbol{\varepsilon}}\frac{Q_1Q_2}{r^2} = \frac{1}{4\pi\boldsymbol{\varepsilon_r\varepsilon_0}}\frac{Q_1Q_2}{r^2} = \left(\frac{1}{4\pi\boldsymbol{\varepsilon_0}}\right)\frac{Q_1Q_2}{\boldsymbol{\varepsilon_r}r^2} = \mathbf{\frac{kQ_1Q_2}{\varepsilon_r r^2}}$$

For ordinary air at sea level, ε_r is almost exactly one (about 1.0006), and we shall treat the force between two charges in air as identical to the force between the same charges in a vacuum, so that no correction is needed when laboratory problems specify "in air". For paraffin, ε_r is about 2.2; for glass, between 4 and 10 depending on the kind; for distilled water, about 80; for the metals, the concept does not apply because free charges inside a conductor simply screen the field and reduce the force to zero. We shall return to this case later.

Coulomb force has a direction as well as a magnitude

The scalar form of Coulomb's law gives the magnitude of the force, but every force has a direction, and the direction of the Coulomb force is fixed by the geometry of the two charges. Each charge feels a force directed along the line joining it to the other charge. If the two charges have the same sign (both positive or both negative), the force on each is directed away from the other, and the two charges repel. If the two charges have opposite signs, the force on each is directed towards the other, and they attract.

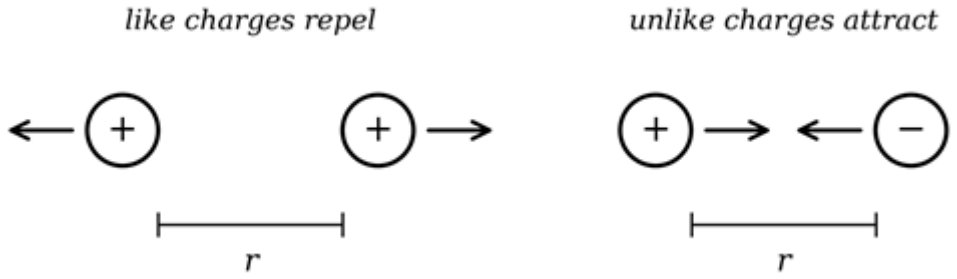


Figure: *Like charges repel along the line joining their centres; unlike charges attract along the same line. In both cases the two forces on the pair are equal in magnitude and opposite in direction, in obedience to Newton's third law. The common separation between the two centres is* r.

The force that charge 1 exerts on charge 2, and the force that charge 2 exerts on charge 1, are equal in magnitude and opposite in direction: this is simply Newton's third law, applied to a pair of charged bodies. Algebraically, writing F_{12} for the force on body 1 due to body 2 and F_{21} for the force on body 2 due to body 1, we have:

$$\mathbf{F_{12} = -F_{21}},$$

and these two opposite vectors share the same line of action, which is the line joining the centres of the two charges. The sign of Q_1Q_2 in the formula encodes the rest: when the product is positive (two like charges) the force is repulsive, and when the product is negative (two unlike charges) the force is attractive. In a problem in which the directions matter, the practical procedure is to compute the magnitude of the force using the magnitudes of the charges, then add the direction by sketching the geometry and applying the like-repels, unlike-attracts rule.

Both pieces of information about the Coulomb force, its size and its direction, can be wrapped into a single **vector equation.** Introduce a unit vector r̂, a vector of length one whose only purpose is to mark the direction along the line joining the two charges, pointing from the first charge to the second. The force on the second charge due to the first can then be written in one line as

$$F = k\frac{Q_1Q_2}{r^2}\,\hat{r}$$

Read the equation in two halves. The fraction $k\frac{Q_1Q_2}{r^2}$ carries the size of the force, exactly as it did in the scalar form, with a sign attached to it that depends on the signs of the two charges. The unit vector r̂ carries the direction. When the product Q_1Q_2 is positive, the two charges sharing the same sign, the equation hands back a vector pointing along r̂, away from the first charge, and the force is the repulsion the rule of signs has always predicted. When the product is negative, the two signs opposite, the equation hands back a vector pointing the other way, along $-\hat{r}$, back towards the first charge, and the force is an attraction. The geometry of attraction and repulsion is no longer a separate rule applied by hand after the calculation; it lives inside the formula itself, governed by the sign of the product of the charges.

The force on the first charge due to the second is, by Newton's third law already noted above, the same vector reversed in direction:

$$F_{12} = -F_{21}$$

Either form of Coulomb's law, the scalar one with directions decided by the rule of signs, or the vector one with directions hidden inside the symbols, is correct. The two are not in competition; they are two ways of saying the same thing.

The principle of superposition

Coulomb's law concerns a pair of charges. Real problems concern collections of three, four, or many more charges. To find the total force on one of the charges produced by all the others, we appeal to a separate experimental fact, called the **principle of superposition.** Stated simply, the principle says that:

The total electric force on a particular charge is the vector sum of the individual Coulomb forces that each of the other charges, taken one at a time, would exert on it.

The principle is not an automatic consequence of Coulomb's law. It is a separate experimental discovery, although a deep one: it amounts to the statement that the Coulomb force between two charges does not depend on the presence or absence of any third charge. The two charges do their own business with each other regardless of the company they keep, and the company's contributions add up by ordinary vector addition. If the principle had failed in some way (if a third charge had caused the force between the first two to shrink, for instance), all of the chapter that follows this paragraph would have to be rewritten.

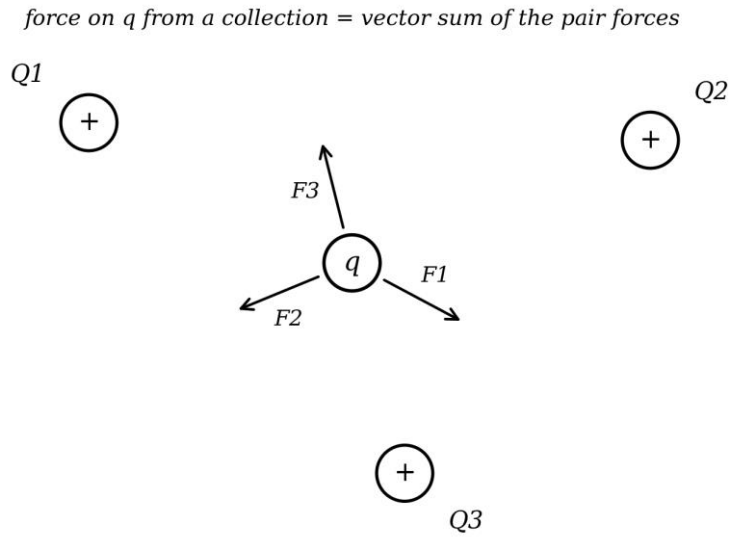


Figure: ***Superposition.*** *Three source charges Q_1, Q_2, and Q_3 each exert a Coulomb force on the target charge q. The total force on q is the vector sum of the three individual forces F_1, F_2, and F_3. Each component force is computed using Coulomb's law as if the other two were not there.*

Symbolically, if q is the charge on which we want the force and Q_1, Q_2, Q_3, and so on are the charges that act on it, then the total force on q is:

$$F_{total} = F_1 + F_2 + F_3 + \cdots,$$

where each F_i is the Coulomb force that the source Q_i alone would exert on q, computed by Coulomb's law and directed along the line joining Q_i to q. The addition is vector addition, with magnitudes and directions both honoured. In problems with symmetric arrangements (charges at the corners of a square, an equilateral triangle, or a regular ring), this addition often collapses by symmetry to a single non-zero component along an axis of symmetry; in less symmetric arrangements, the addition is straightforward but tedious, as we will see later.

In the actual numerical practice of these problems, the vector sum is almost always carried out by components. Choose a pair of axes, the natural pair being the horizontal x and the vertical y, resolve each individual Coulomb force F_i into its components along those axes, and add the components along each axis separately:

$$\sum F_x = F_{1x} + F_{2x} + F_{3x} + \ldots$$
$$\sum F_y = F_{1y} + F_{2y} + F_{3y} + \ldots$$

The two component-sums then combine into a single resultant vector whose magnitude is given by,

$$F = \sqrt{(\sum F_x)^2 + (\sum F_y)^2}$$

and whose direction makes an angle

$$\theta = \tan^{-1}\left(\frac{\sum F_y}{\sum F_x}\right)$$

with the chosen x axis. This four-step procedure, resolve, sum along each axis, combine into a single magnitude, then find the angle, is the procedure every examination-grade worked example follows, and the one that should be reached for the moment three or more charges enter a question.

The twin of Newton's law of gravitation

A reader of our previous book, Foundations of Physics, will feel a strong sense of recognition on meeting Coulomb's law for the first time. The reason is that this is not the first inverse-square law of nature you have met. Newton's universal law of gravitation, which described the force between two masses m_1 and m_2 separated by a distance r, reads:

$$F = \frac{Gm_1m_2}{r^2},$$

where G is the gravitational constant. Set this beside Coulomb's law,

$$F = \frac{kQ_1Q_2}{r^2},$$

and the two structures are identical: a constant of nature, multiplied by the product of two source quantities (masses for Newton, charges for Coulomb), divided by the square of the distance between them. The two laws are twins, and the machinery you will build for one of them transfers, with almost no change of notation, to the other. Gravitational fields, gravitational potentials, gravitational potential energies, and the principle of superposition for gravity all reappear in this chapter under electric names, and the derivations look the same on both sides.

Two important disanalogies must be flagged, however, because every examination paper that touches Coulomb's law tests them.

The first is that *charge has two signs, positive and negative, whereas mass has only one.* Gravity is therefore always attractive, while the electric force is sometimes attractive and sometimes repulsive depending on the signs of the two charges.

The second is that *the electric force is enormously stronger than gravity at the atomic scale.* For an electron and a proton separated by one angstrom, the ratio of the electric force to the gravitational force between them is about 10^{39}, a number so large that gravity is utterly negligible in every atomic and molecular process. The reason a rubbed comb can lift a piece of paper against the entire gravitational pull of the Earth is hidden in that one ratio: at the molecular scale, electric forces are not just stronger than gravity, they are **very much stronger.**

A formula in a frame is one thing; a formula that has computed a numerical answer is another, and is usually the one that sticks. Two worked examples now exercise Coulomb's law on a pair of perfectly ordinary numbers, and then on a pair of numbers so small that the result changes the way you think about both gravity and atoms.

BINDER Example 3

Two small spheres carry charges of $+10\text{nC}$ and $+10\text{nC}$ respectively, and their centres are 5.0cm apart in air. Calculate the magnitude of the electric force on each sphere. Take $k = 9.0 \times 10^9\ \text{Nm}^2/\text{C}^2$.

Solution

Both charges are positive, so the force on each is repulsive and directed along the line joining their centres, away from the other sphere. The magnitudes are equal by Newton's third law, so we need only compute one of them. The magnitude of the Coulomb force between the two spheres in air is given directly by Coulomb's law:

$$F = \frac{kQ_1Q_2}{r^2}$$

Substitute the given values, with every quantity expressed in SI units:

$$F = \frac{(9.0 \times 10^9\ \text{Nm}^2/\text{C}^2)(10 \times 10^{-9}\ \text{C})(10 \times 10^{-9}\ \text{C})}{(0.050\ \text{m})^2} = 3.6 \times 10^{-4}\ \text{N}$$

The magnitude of the force on each sphere is therefore $3.6 \times 10^{-4}\ \text{N}$, or about 0.36mN.

Making Sense of the Answer: *One third of a millinewton sounds like a small force, and on human scales it is. To feel it in your hand you would have to put the equivalent of about a thirty-fifth of a gram on a sensitive kitchen balance, which is below the resolution of most balances available in a school laboratory. Yet ten nanocoulombs is a tiny amount of charge (it is* $\frac{1}{10^{11}}$ *of a coulomb), and the two spheres were five centimetres apart, which is far enough that the inverse-square law has already cut the force back hard. The conclusion to draw is that even very small charges, in a small piece of ordinary laboratory apparatus, produce forces that a precise instrument like a torsion balance can read clearly.*

Think Like a Physicist: *Whenever you read off the answer to a Coulomb-force calculation, look at the three powers of ten in the original data and check that they have arrived at the answer in the right combination. In this problem we had two factors of* 10^{-9} *from the charges in coulombs, one factor of* 10^9 *from the constant k, and two factors of* 100 *from converting centimetres to metres in the denominator. Add up the exponents and you get* -4*: the answer has to be of order* 10^{-4} *newtons. A rough exponent check like this catches most arithmetic slips long before the final number arrives, and it is the single fastest way to know whether you have the order of magnitude right.*

REAL Example 4

In the laboratory of Ilboru Secondary School, two small identical pith balls of mass $m = 0.10\text{g}$ each are hung from a common point by two light insulating threads of length $L = 50\text{cm}$. Each ball is given the same positive charge Q, and the two threads settle in equilibrium at an angle of $\theta = 8°$ from the vertical. Calculate the charge Q on each ball. Take $g = 9.8\text{m/s}^2$ and $k = 9.0 \times 10^9\ \text{Nm}^2/\text{C}^2$.

Solution

Each ball is in equilibrium under three forces: its weight mg acting straight down, the tension T in its supporting thread acting along the thread back to the common pivot, and the Coulomb repulsion F from the other ball, acting horizontally outward. We start with a picture of the geometry and the forces, then write down the two equilibrium equations (one for the horizontal direction, one for the vertical direction), and finally substitute Coulomb's law for F.

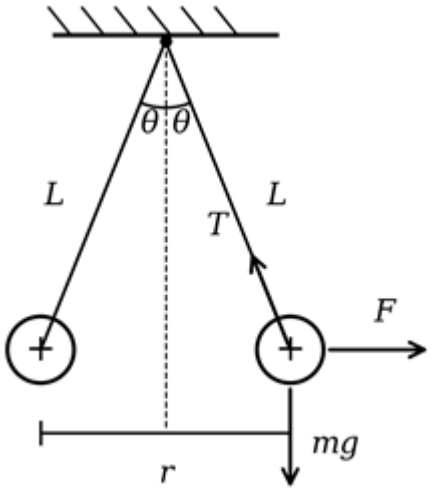


Figure: ***Two pith balls of mass*** **m** ***hang from a common point by threads of length*** **L,** ***each carrying charge*** **Q.** *In equilibrium each thread makes an angle* θ *with the vertical, the centre-to-centre separation is* $r = 2L\sin\theta$*, and each ball feels three forces: its weight* mg*, the thread tension* T*, and the Coulomb repulsion* F.

Resolving the forces on (say) the right-hand ball, the vertical equilibrium condition reads:

$$T\cos\theta = mg,$$

and the horizontal equilibrium condition reads:

$$T\sin\theta = F.$$

Dividing the horizontal equation by the vertical one, T cancels and we are left with:

$$F = mg\tan\theta$$

The geometry fixes the centre-to-centre separation of the two balls. Each ball sits a horizontal distance $L\sin\theta$ from the vertical line through the pivot, so the two balls are separated by:

$$r = L\sin\theta + L\sin\theta = 2L\sin\theta$$

Now write Coulomb's law for the repulsion between the two balls, each carrying the same charge Q:

$$F = \frac{kQ^2}{r^2}$$

Equating the two expressions for F and making Q^2 the subject of the formula:

$$mg\tan\theta = \frac{kQ^2}{r^2}$$

$$Q^2 = \frac{mgr^2\tan\theta}{k} = \frac{mg(2L\sin\theta)^2\tan\theta}{k}$$

Substitute the given values:

$$Q^2 = \frac{(1.0\times10^{-4}\ \text{kg})(9.8\text{m/s}^2)\times(2(0.50\text{m})\sin8°)^2\times\tan8°}{9.0\times10^9\ \text{Nm}^2/\text{C}^2} = 2.97\times10^{-16}\ \text{C}^2$$

$$Q = \sqrt{2.97\times10^{-16}\ \text{C}^2} = 1.72\times10^{-8}\ \text{C}$$

The charge on each ball is therefore 1.72×10^{-8} C or about 17nC.

Making Sense of the Answer: *Seventeen nanocoulombs sits comfortably between the value used in BINDER Example 3 (ten nanocoulombs) and the rough amount of charge a rubbed plastic comb acquires during a few brisk strokes through dry hair. The experiment described here is, in fact, a routine school laboratory demonstration: any teacher with a pair of pith balls, a length of silk thread, and the ability to read a small angle from a protractor can run the experiment in ten minutes and reproduce the calculation we have just done. The pith-ball experiment is the modern descendant of Coulomb's torsion balance, and the formula* $F = mg\tan\theta$ *is the cleanest way ever found to convert a sticky question (how do I weigh an electrostatic force?) into a tidy one (how do I measure an angle?).*

Think Like a Physicist: *When the equilibrium of a charged body involves gravity and a string, divide the horizontal and vertical equations before substituting anything numerical. The tension* T *is usually a quantity nobody cares about, and dividing the two equilibrium equations cancels* T *immediately, replacing two unknowns with one. The same trick applies*

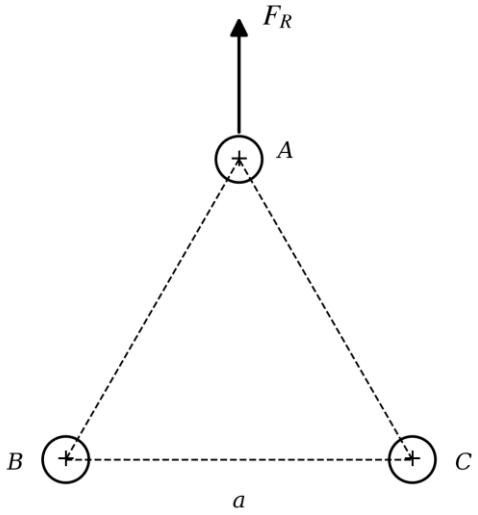
whenever a string holds up a charged or massed object at an angle, and you will see it again in the lightning-rod problems and the suspended oil-drop problems later in this chapter.

HOT Example 5

Three identical positive point charges, each of magnitude $2.0\mu C$, are placed at the vertices of an equilateral triangle of side $30cm$ in air. Calculate the magnitude and direction of the resultant Coulomb force on any one of the charges. Take the Coulomb constant to be $9.0 \times 10^9\ Nm^2/C^2$.

Solution

Label the three vertices A, B, and C, and consider the force on the charge at A. By the principle of superposition, this force is the vector sum of the Coulomb forces exerted on the charge at A by the charges at B and at C. Both contributing forces have the same magnitude, because all three charges are equal and B and C are at the same distance from A (the side of the equilateral triangle). The geometry of the problem, with the resultant force on the charge at A indicated, is shown in the figure.



Let Q be the magnitude of each charge and a be the side of the triangle:

$$Q = 2.0\mu C = 2.0 \times 10^{-6}\ C$$
$$a = 30cm = 0.30m$$

The magnitude of the Coulomb force on the charge at A due to a single neighbour is given by Coulomb's law:

$$F = \frac{kQ^2}{a^2}$$

Substitute the values:

$$F = \frac{(9.0 \times 10^9\ Nm^2/C^2)(2.0 \times 10^{-6}\ C)^2}{(0.30m)^2} = 0.40\ N$$

Each of the two contributions on the charge at A has therefore magnitude $F = 0.40N$. The directions of these two contributions point along the lines BA and CA, extended outward past A (repulsion is from like to like). The two vectors are separated by the interior angle of the equilateral triangle at the vertex A, which is $60°$. By the parallelogram rule, the magnitude of their resultant is:

$$F_R = \sqrt{F^2 + F^2 + 2F^2 \cos 60°}$$

Using $\cos 60° = 0.5$, the expression under the square root collapses to $3F^2$:

$$\mathbf{F_R = \sqrt{3F^2} = F\sqrt{3}}$$

Substitute the value of F:

$$F_R = (0.40N)\sqrt{3} = 0.69N$$

The magnitude is $F_R = 0.69N$.

By symmetry, the direction of the resultant force on the charge at A lies along the perpendicular bisector of the side BC, pointing outward through A away from the centroid of the triangle.

Making Sense of the Answer: *The resultant of two equal forces separated by $60°$ is larger than either alone by a factor of $\sqrt{3}$, not by a factor of two. The reason is that the two contributions do not point in the same direction: they pull at sixty degrees apart, so part of each vector cancels the other's perpendicular component, and only a fraction of each survives in the resultant. The same geometric factor of $\sqrt{3}$ appears at every vertex of every equilateral triangle of like charges, no matter what the charge or the side length, and is worth committing to memory. The result also illustrates the principle of*

superposition in its purest form: the force on one charge from a collection is the vector sum of the pair-force contributions, each computed as if the other source were the only one in the universe.

Think Like a Physicist: *Symmetric arrangements of **equal** charges always reduce the vector-addition step to a single geometric factor. For three equal like charges at the vertices of an equilateral triangle, the factor on each vertex is $\sqrt{3}$. For four equal like charges at the corners of a square, the factor on each corner is the sum of three pair forces whose components partially cancel by symmetry. The rule of thumb: sketch the arrangement, identify the symmetry axis through the chosen charge, project each pair force onto that axis, and the magnitudes parallel to the axis add while the components perpendicular to it cancel. The arithmetic that remains is one Coulomb force times a geometric factor that depends only on how many charges are in the arrangement and where they sit.*

HOT Example 6

In a hydrogen atom, an electron orbits a proton at an average distance of about 5.3×10^{-11} m. The electron has charge $-e$ and mass $m_e = 9.11 \times 10^{-31}$ kg; the proton has charge $+e$ and mass $m_p = 1.67 \times 10^{-27}$ kg; the elementary charge is $e = 1.6 \times 10^{-19}$ C. Calculate (a) the magnitude of the electric force between the electron and the proton, (b) the magnitude of the gravitational force between them, and (c) the ratio of the two forces. Take $G = 6.67 \times 10^{-11}\ \mathrm{Nm^2/kg^2}$.

Solution

(a) The electric force is attractive (the charges are opposite), and its magnitude follows directly from Coulomb's law:

$$F_E = \frac{kQ_1Q_2}{r^2} = \frac{ke^2}{r^2}$$

Substitute the given values:

$$F_E = (9.0 \times 10^9 \mathrm{Nm^2/C^2}) \times \frac{(1.6 \times 10^{-19}\ \mathrm{C})^2}{(5.3 \times 10^{-11}\ \mathrm{m})^2} = 8.2 \times 10^{-8}\ \mathrm{N}$$

The electric force has magnitude $F_E = 8.2 \times 10^{-8}$ N, directed from the electron towards the proton.
(b) The gravitational force is also attractive, and its magnitude follows from Newton's law of gravitation:

$$F_G = \frac{Gm_em_p}{r^2}$$

Substituting:

$$F_G = (6.67 \times 10^{-11} \mathrm{Nm^2/kg^2}) \times \frac{(9.11 \times 10^{-31}\ \mathrm{kg})(1.67 \times 10^{-27}\ \mathrm{kg})}{(5.3 \times 10^{-11}\ \mathrm{m})^2} = 3.6 \times 10^{-47}\ \mathrm{N}$$

(c) The ratio of the electric force to the gravitational force is:

$$\frac{F_E}{F_G} = \frac{8.2 \times 10^{-8}\ \mathrm{N}}{3.6 \times 10^{-47}\ \mathrm{N}} = 2.3 \times 10^{39}$$

The electric force between the electron and the proton in a hydrogen atom is about 2.3×10^{39} times stronger than the gravitational force between them.

Making Sense of the Answer: *The number 10^{39} has a way of disappearing into the page when written down. Make it concrete by saying it aloud: a trillion trillion trillion billion. The electric force pulling the electron onto the proton in a hydrogen atom is a trillion-trillion-trillion-billion times stronger than the gravitational pull between them. If gravity had been the binding force of the atom, the orbit of the electron would have to be unimaginably slow and unimaginably wide to be stable, and the universe of chemistry that builds every plant, animal, and human being you have ever met would simply not exist. The electron clings to the proton because of electricity, and gravity is, on that scale, a rumour. Yet gravity, summed over every atom of an entire planet, produces the force that holds your shoes to the floor of the classroom. The two forces operate at very different scales and very different strengths.*

Think Like a Physicist: *The ratio $\frac{F_E}{F_G}$ computed in this example is independent of the separation r. Both forces fall off as $\frac{1}{r^2}$, so the r^2 in the denominator cancels out, and the ratio depends only on the constants k, G, the two masses, and the elementary charge e. This is a powerful style of argument: whenever two inverse-square forces are compared, the geometry drops out and only the source quantities and the natural constants remain.*

Three more worked examples now lie behind us, and Coulomb's law has done what every great formula in physics does: it has taken a difference of forty orders of magnitude and packed it into a number you can multiply on a single line. The next section asks what would happen if we removed the second charge from these problems and asked, instead, what the first charge alone is doing to the empty space around it. That question takes us out of force-on-a-charge and into the idea of the electric field, which the rest of this chapter is going to live inside.

ELECTRIC FIELD OF A POINT CHARGE

Coulomb's law tells us the force between two charges. It is, in a sense, a complete description of electrostatics: every phenomenon we shall meet from here to the end of the chapter can be derived from it, in principle, by laborious vector addition over many charges. In practice, however, that vector addition is so tedious that a second, more economical idea has, for two

hundred years, been the way working physicists think about electric phenomena. The idea is to forget the second charge for a moment and ask what the first charge, sitting alone, is doing to the empty space around it. The answer, in one sentence, is that *the first charge sets up an electric field everywhere around it, and any other charge that ventures into that field feels a force*. The job of this section is to make the idea of the electric field precise, derive its form for the simplest source, learn to draw it as a picture, and apply it to the most important geometry of this entire chapter: the uniform field between two parallel plates.

From force to field

Imagine that a single source charge Q has been fixed in place somewhere in space, and we want to know what it is doing at a particular point P nearby. The way to find out is to place a small positive test charge q at P and measure the force it feels. The force F on the test charge is, by Coulomb's law, proportional to q: doubling the test charge doubles the force, halving it halves the force, and so on. But the ratio of force to test charge, $\frac{F}{q}$, is independent of how much test charge we used. The ratio depends only on the source Q and on where the test point P is located. We give that ratio its own name and its own symbol.

The **electric field** at a point P is defined as *the electric force per unit positive charge experienced by a small positive test charge placed at that point.* Mathematically, if F is the force on a small positive test charge q at P, the electric field at P is:

$$\mathbf{E} = \frac{\mathbf{F}}{\mathbf{q}}$$

The electric field is a vector: it has the same direction as the force on a positive test charge, and it has the same direction as F divided by the positive scalar q. The SI unit of electric field is therefore the unit of force divided by the unit of charge, i.e., the newton per coulomb: $1\ \mathrm{N/C}$. (A second unit, the volt per metre, will turn out to be identical to the newton per coulomb when we meet the electric potential in a later section, but for now we shall use only the newton per coulomb.)

Two cautions about the test charge q are worth flagging at the start, because students often stumble over them.

- **The test charge must be small.** If we used a test charge so large that its own field disturbed the source Q (by attracting or repelling its charges to new positions), the force we measured would no longer represent the field that Q produces on its own. In practice we imagine q to be vanishingly small.
- **The test charge must be positive by convention**. A negative test charge would feel a force in the opposite direction, and the field defined as F/q would have the wrong sign. The convention "small and positive" gives the field its own definite direction, independent of the test charge we happen to drop into it.

The field of a point charge

Apply the new definition to the simplest possible source: a single point charge Q sitting alone in a vacuum. A small positive test charge q placed at a distance r from Q feels a Coulomb force whose magnitude we already know:

$$F = \frac{kQq}{r^2}$$

Divide this by q as the definition demands:

$$\mathbf{E} = \frac{F}{q} = \frac{\mathbf{kQ}}{\mathbf{r^2}},$$

and we have the field of a point charge, the cleanest formula in electrostatics. The magnitude of the field at distance r from a point charge Q falls as the square of the distance, exactly as the Coulomb force did before it. The direction is fixed by the sign of the source: *if Q is positive the field at P points radially outward from Q (away from the source); if Q is negative the field points radially inward, towards the source.*

In a uniform material medium of relative permittivity ε_r, the constant k is replaced by $\frac{k}{\varepsilon_r}$ (because the medium reduces the electric interaction by the factor ε_r), so the field of a point charge in a medium is:

$$\mathbf{E} = \frac{\mathbf{kQ}}{\boldsymbol{\varepsilon_r}\mathbf{r^2}}$$

directed in the same radial direction as in a vacuum. The medium weakens the field but does not change its direction.

Electric field lines

The electric field at every point in space is a vector, with a magnitude and a direction. To draw such a vector at every point of a region of space, we use a device invented by Michael Faraday and refined by James Clerk Maxwell: the **electric field line.** A **field line** *is an imaginary curve in space whose tangent, at every point along it, gives the direction of the electric field at that point.* A small positive test charge released from rest at any point on a field line would, for an instant, start moving along that line. (Only for an instant, because as soon as the test charge begins to move, it accelerates, gathers speed, and its motion no longer follows the field line. But the tangent rule is precise: the field line shows the direction of the force at every point along it.)

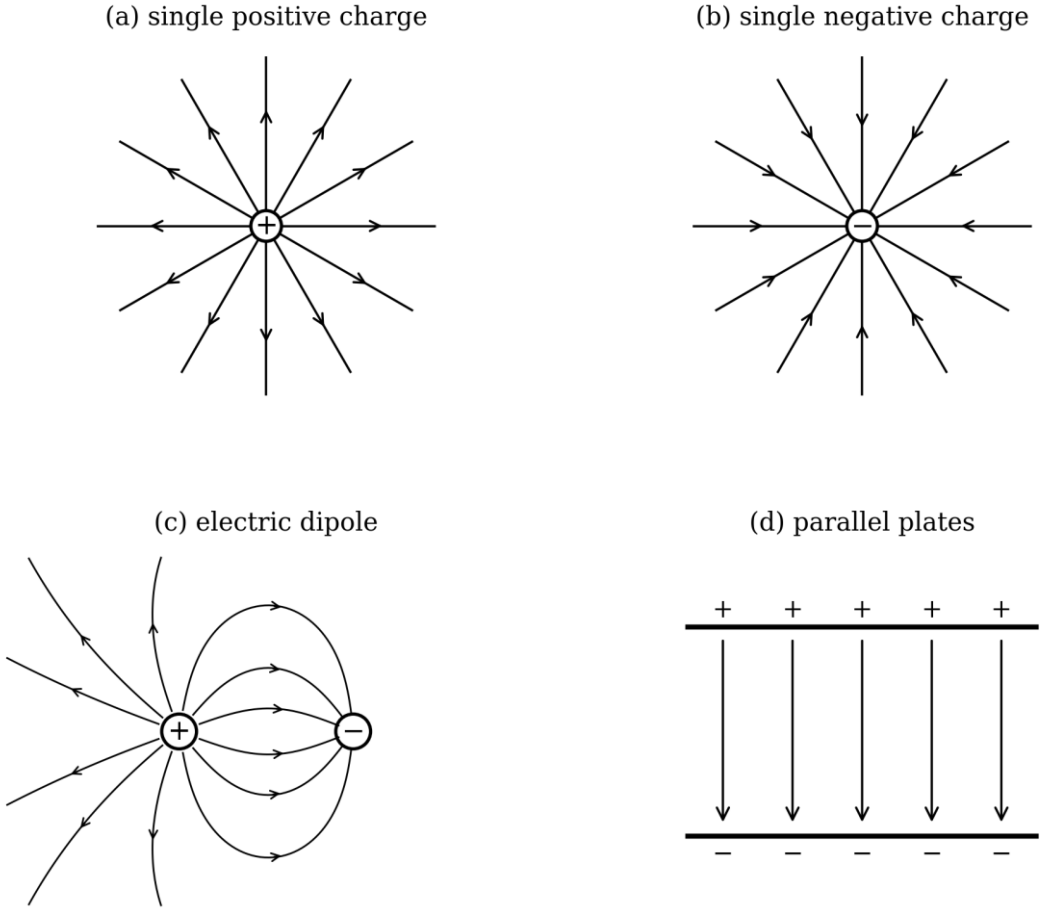


Figure: ***Patterns of electric field lines for four standard charge arrangements****. (a) A single positive point charge: the lines fan out radially in every direction, away from the charge. (b) A single negative point charge: the lines converge radially inward, toward the charge. (c) An electric dipole (one positive charge and one equal negative charge held a fixed distance apart): the lines leave the positive charge, curve through space, and arrive at the negative charge. (d) Two parallel plates carrying opposite uniform charges: the field lines between the plates are straight, parallel, and equally spaced, indicating a uniform field.*

Field-line diagrams are not just pretty pictures. They carry real quantitative information, and they obey six rules that examination questions love to test. The six rules follow from the definition of a field line, from the conservation of charge, and from the way charges arrange themselves on conductors. Understand them so that they are remembered without effort.

Rule 1: Field lines begin on positive charges and end on negative charges

A line cannot start in empty space, because there would be no source for the field along it; nor can it end in empty space, for the same reason. Every line must trace a path from a positive charge to a negative one (or, in the limit, to or from infinity).

Rule 2: Two field lines never cross

If they did, the tangent at the crossing point would have two different directions, and the electric field at that point would be undefined.

Rule 3: The lines are denser (closer together) where the field is stronger, and sparser (further apart) where the field is weaker

The density of lines through any small area perpendicular to the field is proportional to the magnitude of E at that area.

Rule 4: A field line meets the surface of a conductor in static equilibrium at right angles

The reason is that the surface of a static conductor is an equipotential surface, and any component of E parallel to it would drive a current along the surface, in contradiction with static equilibrium.

Rule 5: In electrostatics, field lines never form closed loops

A closed loop would mean that a small positive test charge taken around the loop would return to its starting point with a net gain in kinetic energy, in violation of the conservation of energy. (This rule breaks for the magnetic field of a steady current, where closed loops are the rule rather than the exception. Stay tuned for that surprise in a third book in the series.)

Rule 6: The number of lines drawn entering or leaving a charge is, by convention, proportional to the magnitude of the charge

A charge of $+2Q$ is drawn with twice as many lines leaving it as a charge of $+Q$.

A field-line pattern can be made visible in the laboratory with a simple trick due to Faraday himself. Sprinkle small grass-seed or iron-filing particles on a horizontal sheet of paper, then place two charged objects on the paper a few centimetres apart. Each tiny particle, finding itself in the field, polarises (its electrons shift slightly) and becomes a small electric dipole that aligns itself with the local field direction. After a few seconds, the particles trace out the field-line pattern as if it had been drawn by an invisible hand. The figure shows the patterns the experiment reveals for four arrangements that are worth recognising on sight.

The uniform field between two parallel plates

Among the four field-line patterns of the figure in the previous page, one stands apart for both its mathematical simplicity and its industrial importance: the field between two parallel plates carrying opposite uniform charges. The lines between the plates are straight, parallel, and evenly spaced. By rule 3, an evenly spaced bundle of parallel lines means a field of constant magnitude everywhere in the region they fill. The field between the plates is, in a word, uniform.

A **uniform electric field** *is one whose magnitude and direction are the same at every point of the region of space it fills.* In the gap between two large parallel plates carrying opposite uniform charges, the field is, to an excellent approximation, uniform everywhere away from the plates' edges.

How big is that field? The answer is one of the most useful formulas in the chapter, and we shall derive it properly when Gauss's law arrives in the next section. For the moment, we shall state it and use it. If the two plates are separated by a **perpendicular distance d** and the **potential difference** between them is **V,** then the magnitude of the **uniform electric field** between the plates is:

$$\mathbf{E = \frac{V}{d}},$$

directed from the positive plate to the negative plate. The derivation of this result will follow naturally from the definition of potential difference we shall meet shortly, and from the fact that E is constant in the gap. For now, recognise the formula, know the units (volts per metre, which is the same as newtons per coulomb), and notice three things about it.

This result has several important implications.

First, the field depends only on the ratio $\frac{V}{d}$, not on either quantity separately. A potential difference of 100V across a gap of 1cm produces the same E as 1000V across 10cm.

Second, the formula assumes the plates are large enough, and their separation small enough, that the field between them is genuinely uniform. Near the edges of real plates the field bulges outward and the lines are no longer parallel, an effect called fringing that we shall ignore throughout this chapter except in a single advanced remark much later.

Third, $E = V/d$ is the bridge between two ways of describing what is between the plates. Tell someone the plate separation and the voltage across the plates and they can compute the field. Tell them the field and the distance and they can compute the voltage. The two pictures, field-as-vector and voltage-as-scalar, will turn out to describe the same physics from two complementary angles, and the formula $E = V/d$ is the cleanest example of how the two angles meet.

Two worked examples now arrive to check that we can do something with this new field besides admire it from a distance.

BINDER Example 7

A point charge of $+2.0\mu\text{C}$ is placed inside a tank of paraffin oil, whose relative permittivity 5.0. Calculate the magnitude of the electric field at a point P that lies inside the oil at a distance of 10cm from the charge. Take $k = 9.0 \times 10^9\ \text{Nm}^2/\text{C}^2$.

Solution

The field of a point charge inside a dielectric medium of relative permittivity ε_r is the vacuum field reduced by the factor ε_r:

$$E = \frac{kQ}{\varepsilon_r r^2}$$

Substitute the given values, all expressed in SI units:

$$E = \frac{(9.0 \times 10^9\ \text{Nm}^2/\text{C}^2)(2.0 \times 10^{-6}\ \text{C})}{(5.0)(0.10\text{m})^2} = 3.6 \times 10^5\ \text{N/C}$$

The magnitude of the electric field at P is therefore $3.6 \times 10^5\ \text{N/C}$, directed radially outward from the positive charge.

Making Sense of the Answer: *Three hundred sixty thousand newtons per coulomb is, by the standards of school physics, a strong electric field, larger than the field just below a thundercloud about to discharge, and not far below the breakdown field of dry air. A two-microcoulomb charge is the kind of charge a rubbed plastic rod can hold for a few seconds in dry weather, so the field we have just computed is the field a Form Five student could comfortably produce on a laboratory bench. The factor of five in the denominator (the relative permittivity of paraffin) reduces the field by exactly that factor compared*

with the same charge in air: had the medium been air rather than paraffin, the answer would have been $1.8 \times 10^6\ N/C$*, which is dangerously close to the dry-air breakdown limit we shall meet later.*

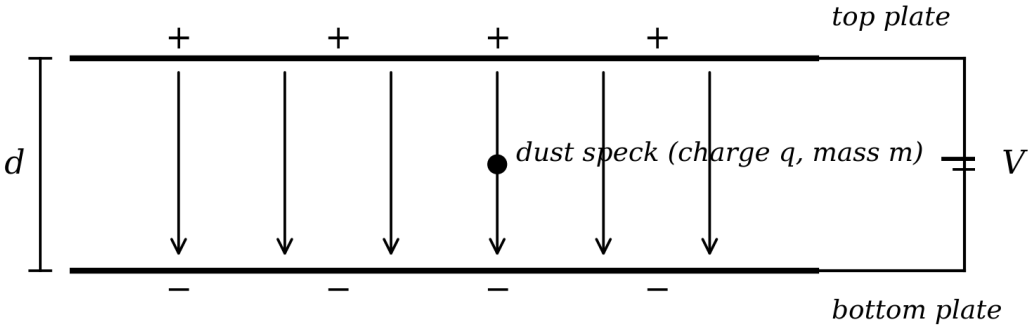
Think Like a Physicist: *Whenever a problem asks for an electric field in a medium rather than in vacuum, the only modification to the vacuum formula is the appearance of the relative permittivity* ε_r *in the denominator. No re-derivation is needed; just divide the vacuum result by* ε_r*. The same scaling applies to the Coulomb force between two charges in the same medium, to the field of any collection of point charges in the medium, and to the capacitance of a parallel-plate capacitor with the same dielectric between its plates. One factor of* ε_r *handles them all.*

REAL Example 8

In a school physics laboratory, two horizontal metal plates are mounted $\mathrm{d} = 4.0\mathrm{cm}$ apart with the upper plate connected to the positive terminal of a $\mathrm{V} = 200\mathrm{V}$ battery and the lower plate connected to the negative terminal. A small spherical dust speck of mass $\mathrm{m} = 5.0\mu\mathrm{g}$ carries a charge of $\mathrm{q} = +1.0\mathrm{nC}$ and finds itself momentarily floating between the plates. Calculate (a) the magnitude of the electric field between the plates, (b) the magnitude of the electric force on the dust speck, and (c) the ratio of the electric force to the weight of the speck. From (c), decide whether the speck will rise, fall, or float. Take $\mathrm{g} = 9.8\mathrm{m/s^2}$.

Solution

Consider the following diagram:



a charged parallel-plate capacitor with a small dust speck between the plates

(a) The plates are large enough and close enough together that the field between them is uniform, with magnitude given by the parallel-plate formula:

$$\mathrm{E} = \frac{\mathrm{V}}{\mathrm{d}}$$

Substitute:

$$\mathrm{E} = \frac{200\mathrm{V}}{0.040\mathrm{m}} = 5.0 \times 10^3\ \mathrm{V/m}$$

The field between the plates has magnitude $\mathrm{E} = 5.0 \times 10^3\ \mathrm{V/m}$ (equivalent to $5.0 \times 10^3\ \mathrm{N/C}$), directed downward from the positive top plate to the negative bottom plate.

(b) The electric force on a charge q sitting in a field E follows from the definition $\mathrm{E} = \mathrm{F/q}$, rearranged for the force:

$$\mathrm{F} = \mathrm{qE}$$

Substitute:

$$\mathrm{F} = (1.0 \times 10^{-9}\ \mathrm{C})(5.0 \times 10^3\ \mathrm{N/C}) = 5.0 \times 10^{-6}\ \mathrm{N}$$

The electric force on the dust speck has magnitude $\mathrm{F} = 5.0 \times 10^{-6}\ \mathrm{N}$, equal to $5.0\mu\mathrm{N}$, directed downward (because the field is downward and the charge is positive).

(c) The weight of the speck is:

$$\mathrm{W} = \mathrm{mg} = (5.0 \times 10^{-9}\ \mathrm{kg})(9.8\mathrm{m/s^2}) = 4.9 \times 10^{-8}\ \mathrm{N}$$

And the ratio of the electric force to the weight is:

$$\frac{\mathrm{F}}{\mathrm{W}} = \frac{5.0 \times 10^{-6}\ \mathrm{N}}{4.9 \times 10^{-8}\ \mathrm{N}} = 1.02 \times 10^2\ = 102$$

The electric force on the speck is about a hundred times its weight, and it points downward in the same direction as the weight. Both forces push the speck downward, so the speck will accelerate downward, not float. It falls to the bottom plate within a fraction of a second.

Making Sense of the Answer: *The answer to part (c) is sometimes surprising on first reading. A microscopic dust speck barely visible to the eye weighs almost nothing, yet a hundred times that almost-nothing is still very, very little. The genuine surprise is the ratio of one hundred: an electric field of just five thousand newtons per coulomb (which a Form Five student*

can produce with two metal plates, a $200V$ battery, and forty millimetres of air) exerts on a charged speck a force a hundred times stronger than gravity. To make the speck float in equilibrium against gravity, the experimenter would have to reduce its charge by a factor of one hundred (down to $10pC$), or weaken the field by the same factor (by turning down the battery to $2V$, or moving the plates four metres apart), or arrange for the field to point upward by reversing the polarity of the battery.

Think Like a Physicist: *Every uniform-field problem with a small charged object between two parallel plates collapses to the same three equations: $E = V/d$ for the field, $F = qE$ for the electric force, and $W = mg$ for the weight. Once those three numbers are in hand, the question of whether the object rises, falls, or floats reduces to comparing F and W, and to deciding which way F points (which depends on the sign of the charge and the orientation of the plates). The same machinery, with the same three equations, will reappear in the motion of charged particles in fields later in this chapter, in the discussion of the Millikan oil-drop experiment, and in the analysis of charged droplet steering in inkjet printers covered near the end of the chapter.*

Two worked examples now lie behind us, and the electric field has earned its keep. We can compute it for a single point charge in any medium, we can draw it as a pattern of lines for any arrangement we can sketch, and we can read off its strength in a parallel-plate gap from a single division. The natural next question is what happens when the charge is no longer a single point sitting on its own, but is smeared along a curve, over a surface, or through a volume. That is exactly the question the next section is about, and the principle of superposition we met in the previous section is going to make the answer beautifully tractable.

ELECTRIC FIELD FROM DISTRIBUTIONS OF CHARGE

A single point charge produces a tidy, radial field that falls as the square of distance from it. Real problems rarely come with a single point charge sitting on its own. They come with two charges at the ends of a thin rod, four charges at the corners of a square, a ring of bracelets-worth of charge, an entire long wire of charge running across a laboratory bench, or, in due course, a charged conducting sphere with all its charge spread across the surface. The job of this section is to take the principle of superposition we met in the previous one and convert it into a tool for finding the field of any reasonable distribution of charges. We will do it three times in three increasingly continuous settings: a finite collection of point charges sitting at the corners of a square; the continuous distribution of charge around a thin ring; and, as a final case that motivates the deep theorem of the next section, an infinite line of charge.

A collection of point charges

The **principle of superposition for electric fields** states that:

The total electric field at a point due to a collection of source charges is the vector sum of the individual electric fields that each source, taken one at a time, would produce at that point.

The principle transfers straight from forces to fields, since dividing the total force on a small positive test charge by the test charge simply divides each source-contribution to the force by the same amount. Algebraically, if a small positive test charge q sitting at point P would feel a force $\mathrm{F_i}$ from the *i*th source charge if that source were alone in the universe, then the total electric field at P produced by a whole collection of source charges acting simultaneously is the vector sum of the individual fields:

$$\mathrm{E} = \mathrm{E_1} + \mathrm{E_2} + \mathrm{E_3} + \cdots$$

Each $\mathrm{E_i}$ is computed by treating the corresponding source as if it were alone, using the point-charge formula for the field and pointing the vector along the line from the source to P (radially outward for a positive source, radially inward for a negative one). The arithmetic of the calculation is always the same: write each contribution as a vector, resolve the vectors into components, add the components, recombine into a single resultant vector. In problems with an obvious symmetry, the addition often collapses dramatically because several pairs of components cancel each other, leaving only one or two non-zero terms.

A continuous distribution: the field on the axis of a ring

When the charges are no longer countable on the fingers of one hand but are smeared continuously over a curve, a surface, or a volume, the sum of individual contributions becomes an integral. The simplest such problem with a clean closed-form answer, is the field on the axis of a uniformly charged thin ring. We have a thin ring of radius R, lying flat in a horizontal plane, and the total charge Q is distributed evenly around the ring. We want the electric field at a point P on the axis of the ring, at a perpendicular distance x from the centre.

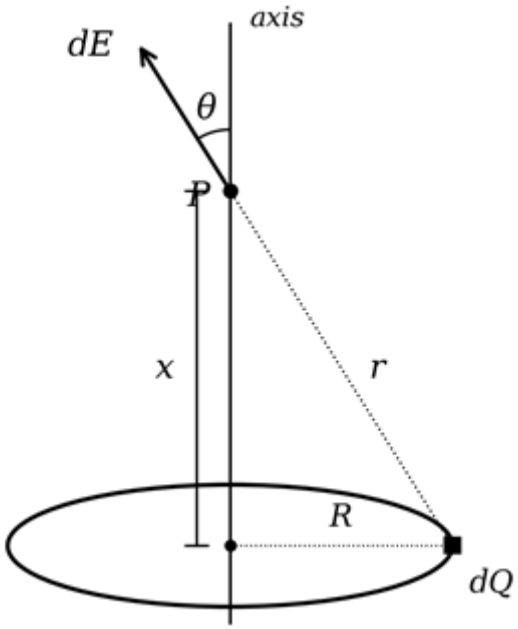


Figure: ***A uniformly charged ring of radius R and total charge Q.*** *The field point P sits on the axis at perpendicular distance x from the ring's centre. A typical charge element dQ on the ring is at distance r from P; the line dQ-to-P makes an angle θ with the axis. The field dE produced at P by the element points along this line, away from the element for a positive dQ.*

Pick any one small piece of the ring, a tiny element of charge dQ. By the point-charge formula, that single element produces a field at P of magnitude:

$$dE = \frac{k\,dQ}{r^2}$$

Where r is the distance from the element to P. From the figure's geometry, that distance is the hypotenuse of a right triangle whose legs are R and x, so:

$$r = \sqrt{R^2 + x^2}$$

The vector dE points along the line from the element to P, so it has two components: one along the axis (pointing away from the ring towards P), and one perpendicular to the axis (pointing in the plane that contains the element, the axis, and P). If we call the angle between dE and the axis θ, then the axial component of dE is $dE\cos\theta$ and the perpendicular component is $dE\sin\theta$.

Here is where the symmetry of the ring does its job. For every charge element on the ring at the position shown in the figure, there is an exactly equal element on the diametrically opposite side of the ring. That opposite element contributes its own dE, equal in magnitude to the first one but pointing in a direction that is the mirror image of the first across the axis. The axial components of the two add together (both point along the axis in the same direction), while the perpendicular components are equal and opposite, and cancel each other exactly. Pair up every element on the ring with its diametrically opposite partner and the result is the same: all perpendicular components cancel in pairs, and only the axial component survives.

So the field at P is purely axial. Its magnitude is the integral of the axial contribution $dE\cos\theta$ over every element on the ring:

$$E = \oint_{\text{ring}} (dE\cos\theta) = \oint_{\text{ring}} \left(\frac{k\,dQ}{r^2}\cos\theta\right)$$

To evaluate this integral, parametrise the ring by the angle φ that a typical element makes around the centre, with φ running from 0 to 2π as we go once around. The arc length of an element subtending an angle dφ is R dφ. If the linear charge density on the ring is $\lambda = \frac{Q}{2\pi R}$, then the element carries:

$$dQ = \lambda R\,d\varphi$$

From the geometry, every element on the ring is at the same distance $r = \sqrt{R^2 + x^2}$ from P, and every element makes the same angle θ with the axis. From the figure, $\cos\theta = \frac{x}{r}$. Substituting these expressions for dQ and $\cos\theta$ into the integrand:

$$E = \int_0^{2\pi} \left(\frac{k(\lambda\,R\,d\varphi)}{r^2} \times \frac{x}{r}\right) = \int_0^{2\pi} \left(\frac{k\lambda R x}{r^3}\right) d\varphi$$

Now look at the integrand: the factors k, λ, R, x, and r are all constants for a given field point P (they do not depend on φ). They pull out of the integral as a common factor, leaving only the integral of dφ:

$$E = \frac{k\lambda Rx}{r^3}\int_0^{2\pi} d\varphi = \frac{k\,\lambda\,R\,x}{r^3} \times 2\pi = \frac{k(2\pi R\lambda)x}{r^3}$$

But $2\pi R\lambda$ is the total charge on the ring (the linear charge density multiplied by the full circumference), and that is just Q. So our expression collapses to:

$$E = \frac{kQx}{r^3}$$

Substituting $r = \sqrt{R^2 + x^2}$ gives the final result:

$$\mathbf{E(x)} = \frac{kQx}{\left(\sqrt{R^2 + x^2}\right)^3} = \frac{\mathbf{kQx}}{(\mathbf{R^2 + x^2})^{3/2}},$$

directed along the axis, away from the ring (for a positive Q). This is the central formula of the subsection: the magnitude of the electric field at a point on the axis of a uniformly charged ring, at perpendicular distance x from the ring's centre.

Two limiting cases of this formula are worth thinking about before we set numbers into it.

First, at the centre of the ring ($x = 0$), the formula gives $E(0) = 0$. The field is zero at the centre. This may, at first sight, seem surprising: there is charge everywhere around the centre, and a great deal of it close to hand. But each element of charge on the ring has a partner on the opposite side, and the two elements pull a test charge at the centre in exactly opposite directions with exactly equal magnitudes. The pulls cancel pair by pair, and the net field at the centre is zero. Symmetry has done what no amount of arithmetic could: it has predicted the answer without computing a single contribution.

Second, very far from the ring ($x \gg R$), the radius R becomes negligible compared with x, and the formula reduces to $E \approx \frac{kQ}{x^2}$, the field of a point charge Q at distance x. This makes perfect sense: from far enough away, the ring looks like a single point with all of its charge piled at the centre, and the field should reduce to the point-charge form.

Where on the axis is the field largest?

The two limits we have just identified are useful but they do not exhaust the geometry of $E(x)$. The field is zero at $x = 0$, it falls back to zero as $x \to \infty$, and it is positive in between. By continuity, the field must reach a maximum somewhere on the axis between the centre and infinity. The location of that maximum is fixed by setting the derivative of $E(x)$ to zero:

$$\frac{dE}{dx} = 0$$

Because $E(x)$ is a quotient of two functions of x, the derivative is computed by the quotient rule. Identify the numerator and the denominator:

$$u(x) = kQx, \quad v(x) = (R^2 + x^2)^{3/2}$$

so that $E(x) = u(x)/v(x)$. The quotient rule reads:

$$\frac{dE}{dx} = \frac{v\,\frac{du}{dx} - u\,\frac{dv}{dx}}{v^2}$$

Compute the two derivatives separately. The numerator derivative is immediate:

$$\frac{du}{dx} = \frac{d}{dx}(kQx) = kQ$$

For the denominator, apply the chain rule. Let $w = R^2 + x^2$, so $v = w^{3/2}$ and $dw/dx = 2x$.

Differentiate $w^{3/2}$ to get $\frac{3}{2}w^{1/2}$, then multiply by $\frac{dw}{dx}$:

$$\frac{dv}{dx} = \frac{3}{2}(R^2 + x^2)^{1/2} \times (2x) = 3x(R^2 + x^2)^{1/2}$$

Substitute these two derivatives into the quotient rule, pulling the common factor kQ out of the numerator:

$$\frac{dE}{dx} = kQ \times \frac{(R^2 + x^2)^{3/2} - 3x^2(R^2 + x^2)^{1/2}}{(R^2 + x^2)^3}$$

Setting this derivative to zero, the denominator is irrelevant (a fraction is zero only when its numerator is), and the numerator must vanish:

$$(R^2 + x^2)^{3/2} = 3x^2(R^2 + x^2)^{1/2}$$

Divide both sides by the common factor $(R^2 + x^2)^{1/2}$ and rearrange:

$$R^2 + x^2 = 3x^2 \implies R^2 = 2x^2$$

Taking the positive square root gives the location:

$$\mathbf{x = \frac{R}{\sqrt{2}}}$$

The field rises from zero at the centre, climbs to its highest value at a distance equal to the radius divided by the square root of two, then falls back to zero in the point-charge limit. The shape of the curve is shown in the following figure.

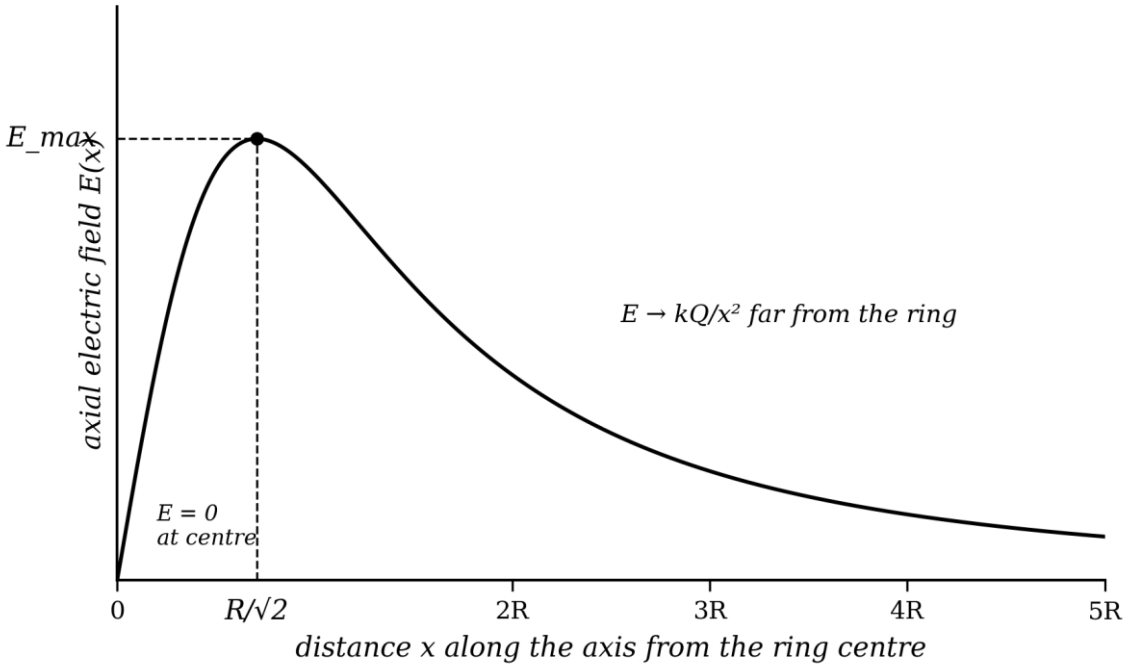


Figure: *The axial electric field of a uniformly charged ring as a function of distance* x *along the axis from the ring centre. The field vanishes at the centre (*$x = 0$*), rises to a maximum at* $x = R/\sqrt{2}$*, and decays as* $\frac{kQ}{x^2}$ *far from the ring, in agreement with the point-charge limit.*

Field of a finite line of charge

After the closed loop of the ring, the most natural continuous distribution to study next is a single straight line of charge. Take a thin straight rod of length 2a, carrying a positive charge spread uniformly along its length with linear charge density λ (we will define it more clearly in the next subsection). Set up the geometry so that the rod lies along the y axis with its midpoint at the origin, and let the field point P sit on the x axis at perpendicular distance x from the midpoint, as the figure shows.

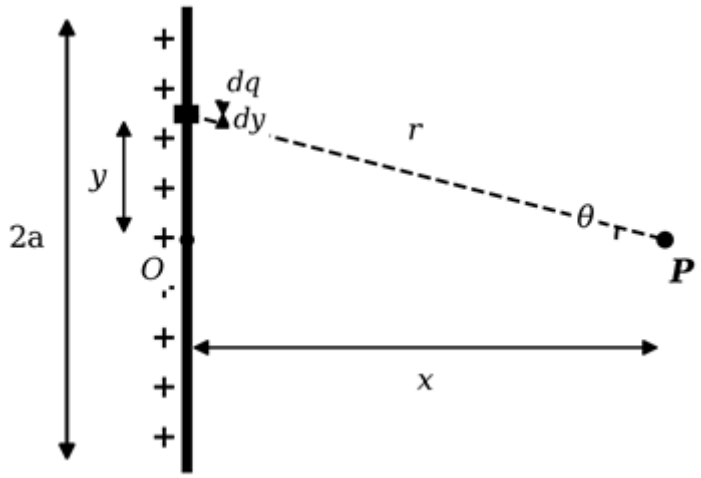


Figure: ***Geometry for the finite line of charge.*** *A thin straight rod of length 2a, with linear charge density λ, runs along the y axis. The field point P sits at perpendicular distance x from the midpoint O. A representative element of length dy at height y above the midpoint carries charge dq = λ dy, sits at distance r from P, and contributes a field dE whose direction is along the line from the element to P.*

By the symmetry of the figure, a charge element at height y above the midpoint and the matching element at height $-y$ below it contribute equal magnitudes of dE at P, but with their y components opposite. The two y components cancel in the sum, and only the components along the x axis survive. The x component of the field due to a single element of charge $dq = \lambda\, dy$ at height y is the Coulomb-law magnitude multiplied by the geometric factor x/r, giving:

$$dE_x = dE\cos\theta = \frac{k\,dq}{r^2} \times \frac{x}{r} = \frac{k\lambda x\, dy}{(x^2 + y^2)^{3/2}}$$

With x held fixed and the variable y running over every element of the rod from $-a$ to a, the total field at P is the integral:

$$E_x = k\lambda x \int_{-a}^{a} \frac{dy}{(x^2 + y^2)^{3/2}}$$

The standard integral evaluates to:

$$\int_{-a}^{a} \frac{dy}{(x^2 + y^2)^{3/2}} = k\lambda x \left[\frac{y}{x^2\sqrt{x^2 + y^2}}\right]_{-a}^{a} = k\lambda x \left(\frac{2a}{x^2\sqrt{x^2 + a^2}}\right)$$

Therefore the field of the finite line of charge at perpendicular distance x from the midpoint is:

$$\mathbf{E_x = \frac{2k\lambda a}{x\sqrt{x^2 + a^2}}}$$

The result behaves correctly at the two natural extremes.

- When the field point sits very close to the rod, x much smaller than a, the square root reduces to a, and the formula collapses to $\mathbf{E = \frac{2k\lambda}{x}}$, the field of a uniformly charged wire treated as effectively **infinite**, which is the case the next subsection takes up directly.
- When the field point sits very far from the rod, x much larger than a, the square root is dominated by x, the rod shrinks to a point of total charge $Q = 2\lambda a$, and the formula collapses to $\mathbf{E = \frac{kQ}{x^2}}$, the Coulomb field of a point charge.

Between these two extremes the full formula gives the field of any actual finite rod.

An infinite line of charge

One last continuous distribution deserves attention before we close this section, because it both gives a beautiful result and exposes the limits of the brute-force summation method we have been using. Consider an infinite, thin straight wire on which positive charge is distributed uniformly along its length.

The **linear charge density** of a charge distribution along a curve is defined as *the amount of electric charge per unit length of the curve*. We give it the symbol λ, and its SI unit is the coulomb per metre, C/m. If a thin straight piece of wire of length ℓ carries a total charge Q spread uniformly along its length, then:

$$\lambda = \frac{Q}{\ell}$$

The question we want to answer is*: what is the electric field at a point P that sits at a perpendicular distance s from such an infinite line of charge?*

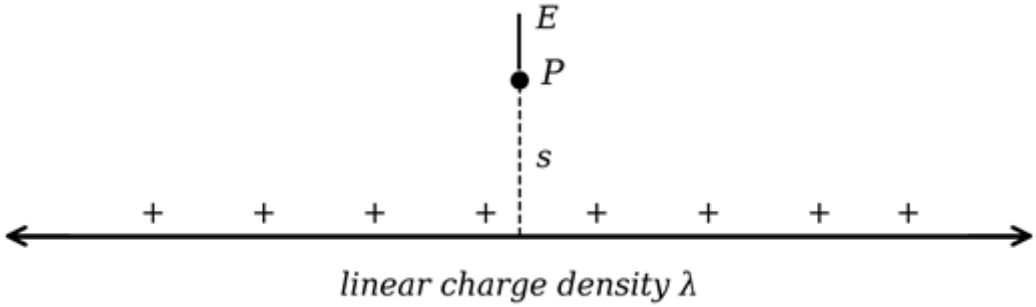


Figure: ***An infinite straight wire carries a uniform linear charge density λ****. The field point P sits at perpendicular distance s from the wire, and by symmetry the field at P points radially outward from the wire, perpendicular to its length.*

Set the perpendicular foot of P as the origin and let y be the coordinate of a point on the wire measured from that foot. A small element of the wire of length dy at position y carries charge,

$$dQ = \lambda dy,$$

and lies at a distance,

$$r = \sqrt{s^2 + y^2}$$

from P. The field dE that this element produces at P has magnitude,

$$dE = \frac{kdQ}{r^2} = \frac{k\lambda dy}{s^2 + y^2},$$

and points along the line from the element to P. Let α be the angle this line makes with the perpendicular (the perpendicular itself is the direction of the field we want). The perpendicular component of dE is $dE\cos\alpha$, and the geometry gives $\cos\alpha = \frac{s}{r}$. By symmetry, the components parallel to the wire from the elements at $+y$ and $-y$ cancel exactly, so only the perpendicular component survives in the integral over the whole wire:

$$E = \int_{-\infty}^{+\infty} (dE\cos\alpha) = \int_{-\infty}^{+\infty} \left(\frac{k\lambda dy}{s^2 + y^2} \times \frac{s}{r}\right)$$

Substituting $r = \sqrt{s^2 + y^2}$ into the denominator and using the symmetry of the integrand about $y = 0$ to fold the integral onto the half-line $y = 0$ to $y = +\infty$ (which doubles its value):

$$\mathbf{E} = 2k\lambda s \int_0^{+\infty} \frac{dy}{(s^2 + y^2)^{3/2}} = 2k\lambda s \left[\frac{y}{s^2\sqrt{s^2 + y^2}}\right]_0^{+\infty} = 2k\lambda s \times \frac{1}{s^2} = \mathbf{\frac{2k\lambda}{s}}$$

Finally, substituting $k = \frac{1}{4\pi\varepsilon_0}$ to write the result in the standard SI form:

$$\mathbf{E} = \mathbf{\frac{\lambda}{2\pi\varepsilon_0 s}},$$

directed radially outward from the wire. Two features of this result deserve underlining.

First, the field falls with the first power of the distance, not the square. The reason is geometric: a line of charge is one-dimensional, so its field spreads out over a circumference (a one-dimensional surface) of length $2\pi s$, rather than over a sphere of area $4\pi s^2$ as a point charge does.

Second, the answer to this problem fell out of an integral that, while standard, took about half a page of calculus to evaluate. It would be nice to have a shortcut, an organising principle that lets us read off the answers to problems with high symmetry without having to set up and evaluate any integrals at all. That shortcut exists, and it is the subject of the next section. The principle is called Gauss's law, and once you have it in hand, you will solve in three lines the problem we have just solved in half a page.

Three worked examples now exercise the machinery just built.

HOT Example 9

Four point charges sit at the corners of a square of side $a = 20\text{cm}$. The two charges at the top corners are each $+Q = +1.0\mu\text{C}$, and the two at the bottom corners are each $-Q = -1.0\mu\text{C}$. Find the magnitude and direction of the electric field at the centre P of the square. Take $k = 9.0 \times 10^9\ \text{Nm}^2/\text{C}^2$.

Solution

Consider the following figure:

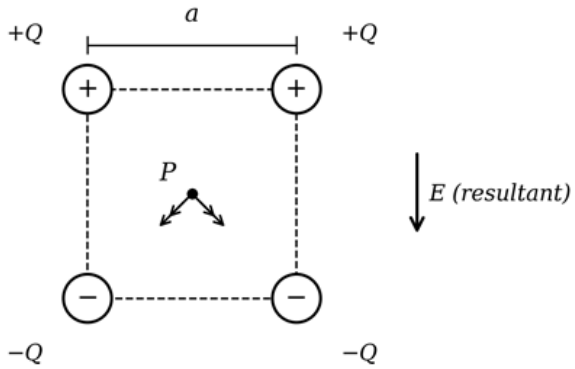


Figure: ***The four-corner arrangement of HOT Example 9****. The two top charges are $+Q$; the two bottom charges are $-Q$. By symmetry, the horizontal components of the four field contributions at the centre cancel in pairs, and the resultant E points straight down from the $+Q$ row to the $-Q$ row.*

The centre of the square is equidistant from all four corners. That distance is half the diagonal of the square:

$$r = \frac{a\sqrt{2}}{2} = \frac{a}{\sqrt{2}}$$

The magnitude of the field at P due to any one of the four charges, taken alone, is given by the point-charge formula:

$$E_1 = \frac{kQ}{r^2} = \frac{kQ}{\frac{a^2}{2}} = \frac{2kQ}{a^2}$$

Each of the four field contributions at P has this magnitude. The directions are fixed by the sign and position of each corner charge. The contribution from each top corner (a positive charge) points away from that corner, i.e., downward and into the square. The contribution from each bottom corner (a negative charge) points towards that corner, i.e., also downward and into the square. Each of the four contributions makes an angle of $45°$ with the vertical, because the line from any corner to the centre runs along the half-diagonal at this angle.

Now resolve each contribution into horizontal and vertical components. The horizontal components are: from the top-left (positive) charge, rightward (positive *x*); from the top-right (positive) charge, leftward (negative *x*); from the bottom-left

(negative) charge, leftward (negative *x*); from the bottom-right (negative) charge, rightward (positive *x*). Two rightward and two leftward, all of equal magnitude: they cancel in pairs, and the horizontal component of the total field is zero.

The vertical components are all downward: from the two top positive charges, the field at P points away (downward into the square); from the two bottom negative charges, the field at P points towards them (also downward). All four vertical components add. Each has magnitude $E_1 \times \cos 45°$, and the total vertical field is the sum of all four:

$$E = 4 \times E_1 \times \cos 45° = 4 \times \frac{2kQ}{a^2} \times \frac{1}{\sqrt{2}} = \frac{8kQ}{a^2\sqrt{2}} = \frac{4\sqrt{2}\,kQ}{a^2}$$

Substitute the given values:

$$E = \frac{(4)\,(1.414)\,(9.0 \times 10^9 Nm^2/C^2\,)\,(1.0 \times 10^{-6} C)}{(0.20m)^2} = 1.27 \times 10^6\ N/C$$

The electric field at the centre of the square has magnitude approximately $1.3 \times 10^6\ N/C$, and it points straight down, from the row of positive charges towards the row of negative charges.

Making Sense of the Answer: *A field of just over one million newtons per coulomb sounds enormous on first reading. It is, in fact, the field of a small commercial Van de Graaff generator: the kind of device that can lift the hair of a child standing on the insulating mat and make a fluorescent tube glow when held near. The four-corner arrangement is the smallest collection of point charges that produces a non-zero, easily computable field at its geometric centre, and the answer* $4\sqrt{2}\ kQ/a^2$ *is the cleanest combination of the elementary constants and the geometry.*

Think Like a Physicist: *Before substituting any numbers into a discrete-charge problem, sketch the arrangement on paper, draw the field vector from every source charge to the field point, and inspect the picture for symmetry. In this problem the four contributions to the horizontal component of the field cancelled in pairs by inspection, and the four vertical components added. Symmetry replaced a four-vector addition with the multiplication of one scalar by four. The same collapse occurs in every regular-polygon arrangement of charges, and recognising it before you start computing is the single largest time-saver in this entire chapter.*

REAL Example 10

In the school physics laboratory, Mr. Akilikubwa sets a thin wire ring of radius $R = 10cm$ on the bench and rubs it carefully with a piece of silk until it carries a uniform positive charge of about $Q = 50nC$. He then asks the class: ***what is the electric field at the centre of the ring?*** Kipanga and Kipute disagree.

Kipanga: *"The field at the centre must be enormous. There is charge everywhere round the ring, only ten centimetres away in every direction. Adding up so many close-by contributions has to give a big answer."*

Kipute: *"Hold on. There is charge everywhere round the ring, but for every bit of charge on one side of the centre there is an equal bit on the opposite side. The two bits pull a small test charge at the centre in exactly opposite directions, and the pulls have exactly the same size because the two bits are the same distance away. They cancel. Every bit cancels its partner. The total field at the centre is zero."*

Kipanga: *"Zero? Surrounded by all that charge, the field is zero? That sounds wrong. Kipute, you are over-relying on symmetry."*

Mr. Akilikubwa: *"Let us settle the matter with the formula we derived a moment ago. The axial field of a uniformly charged ring at distance x from the centre is"*

$$E(x) = \frac{kQx}{(R^2 + x^2)^{3/2}}$$

Mr. Akilikubwa: *"The centre of the ring is the special case $x = 0$. Substitute that value into the formula and read off the answer."*
Kipute carries out the substitution at the chalkboard:

$$E(0) = \frac{kQ \times (0)}{(R^2 + 0)^{3/2}} = \frac{0}{R^3} = 0$$

Kipute: *"Zero. Just as I said. The numerator has an x, and when x is zero the whole thing is zero, no matter how big Q or how small R is."*

Kipanga: *"All right, the formula says zero. But what does the formula know that I do not?"*

Mr. Akilikubwa: *"It knows that the directions of the contributions matter, not just their sizes. A bit of charge on the right of the ring pushes a test charge at the centre to the left. A bit on the left pushes it to the right. The two pushes are equal in size, and they point in opposite directions, so they add to zero, not to twice one of them. Kipanga, your instinct that the answer should be big was based on adding up only the sizes of the contributions, forgetting their directions. The formula folds direction in for you automatically, through the cosine of the angle the contribution makes with the axis. At the centre, the cosine is zero for every element on the ring, and the axial component of each contribution is zero. There is your answer."*

Kipute writes the final answer on the board:

$$\mathbf{E(0) = 0}$$

Mr. Akilikubwa: *"And the lesson: when a charge distribution is symmetric, trust the symmetry. It will save you the trouble of the calculation, and it will give you the answer with no arithmetic at all."*

Making Sense of the Answer: *Kipanga's intuition was perfectly reasonable, and it belongs to a long tradition of students for whom a lot of charge close-by should produce a large field. The tradition is half right: a lot of charge close-by does produce a lot of field, contribution by contribution. The other half is that the contributions are vectors, and vectors add by cancellation as easily as they add by reinforcement. The ring at its centre is the cleanest example of perfect cancellation in all of electrostatics: every contribution has an exact opposite, and the sum collapses to zero. The same principle reappears in capacitor problems, in dipole problems, and in the metallic-cage problems further on in this part of the book.*

Think Like a Physicist: *Whenever a problem asks for the field at the geometric centre of a symmetric distribution of like charges, the answer is almost certainly zero, and you can write it down without setting up any integrals. The same answer holds at the centre of a uniformly charged sphere, at the centre of a uniformly charged spherical shell, at the centre of a circular disc of uniform charge, at the centre of a regular polygon of equal point charges, and at the centre of a uniformly charged cube. Symmetry is the cheapest shortcut in physics: it costs nothing to use, and it replaces an integral with the single word "zero".*

HOT Example 11

A thin ring of radius 10cm carries a uniformly distributed positive charge of 50nC. Find (a) the location along the axis where the electric field is greatest, and (b) the magnitude of the field at that point.

Solution

(a) The axial field of the ring is:

$$\mathrm{E(x)} = \frac{\mathrm{kQx}}{(\mathrm{R^2 + x^2})^{3/2}}$$

The field is largest at the axial distance:

$$\mathrm{x} = \frac{\mathrm{R}}{\sqrt{2}}$$

With $\mathrm{R} = 0.10\text{m}$, this places the maximum at:

$$\mathrm{x} = \frac{0.10\text{m}}{\sqrt{2}} = 0.0707\text{m} = 7.1\text{cm}$$

(b) Substituting $\mathrm{x} = \mathrm{R}/\sqrt{2}$ into the expression for $\mathrm{E(x)}$:

$$\mathrm{E_{max}} = \frac{\mathrm{kQ} \times \dfrac{\mathrm{R}}{\sqrt{2}}}{\left(\mathrm{R^2} + \dfrac{\mathrm{R^2}}{2}\right)^{3/2}} = \frac{\dfrac{\mathrm{kQR}}{\sqrt{2}}}{\dfrac{3\sqrt{3}}{2\sqrt{2}}\mathrm{R^3}} = \frac{2\mathrm{kQ}}{3\sqrt{3}\mathrm{R^2}}$$

Substituting numerical values, with $\mathrm{k} = 9.0 \times 10^9\ \mathrm{Nm^2/C^2}$, $\mathrm{Q} = 50 \times 10^{-9}\ \mathrm{C}$, and $\mathrm{R} = 0.10\text{m}$:

$$\mathrm{E_{max}} = \frac{2\,(9.0 \times 10^9 \mathrm{Nm^2/C^2})\,(50 \times 10^{-9}\mathrm{C})}{3\sqrt{3(0.10\text{m})^2}} = 1.73 \times 10^4\ \mathrm{N/C}$$

The maximum axial field has magnitude approximately $1.7 \times 10^4\ \mathrm{N/C}$, located on the axis at $\mathrm{x} = 7.1\text{cm}$ from the centre of the ring.

Making Sense of the Answer: *A field of seventeen thousand newtons per coulomb is in the normal range for charged objects of laboratory size. The location of the maximum, at* $\mathrm{R}/\sqrt{2}$*, is one of the small geometric facts that physics rewards students for remembering.*

Think Like a Physicist: *When a function vanishes at both ends of an interval and is positive in between, it must have a maximum, and you locate it by setting the derivative to zero.*

Three worked examples now lie behind us, and the superposition principle has done its job on three increasingly continuous geometries. The next section gives us a powerful new tool, Gauss's law, that will reduce the infinite-line problem to three lines instead of a page.

GAUSS'S LAW AND FOUR CANONICAL APPLICATIONS

Last lesson finished with three integrals stacked one on top of the other. A ring of charge took us through a circle. An infinite line took us through the entire real number line. A square arrangement asked for four vector additions and a Pythagorean theorem before it agreed to give up its answer. The work was honest, but it was work. A clever person, watching the same problems for the third time, begins to suspect that there must be a shortcut hiding underneath. There is. Its name is **Gauss's law,** and the rest of this section is the story of how it cuts a class of hard electrostatic problems down to a single line of algebra.

Kipanga, Kipute, and Mr. Akilikubwa are once again on the veranda. Kipanga is balancing a textbook on his head, partly out of mischief and partly because his cousin once told him that physicists develop better posture than other people. Kipute is sketching a sphere on her notebook. Mr. Akilikubwa is, as always, watching them both with the patient amusement of someone who has seen this moment many times before.

Kipute: *Sir, the ring example showed me how to do the integral, but I will not be honest if I claim the calculation felt natural. Is there a way to get the same answer without all the trigonometry?*

Mr. Akilikubwa: *For a ring on its axis, no. The geometry asks for an integral and you must do the integral. But for a charged sphere, an infinite plane, or an infinite line, there is a beautiful counting argument that delivers the answer in one line. The counting argument is called Gauss's law. Before we can state it, however, we need one new word: flux.*

Kipanga: *Flux. That sounds like a disease.*

Mr. Akilikubwa: *It is in fact the Latin word for flow. We will use it to count field lines crossing a surface.*

Electric flux

Imagine a flat surface of area A placed inside a uniform electric field of magnitude E. Suppose first that the surface is held perpendicular to the field, with the field lines piercing it head-on. The number of field lines crossing the surface is proportional to the product EA. That product is called the **electric flux** through the surface, written with the Greek capital letter Φ (phi).

$$\boldsymbol{\Phi = EA}$$

The electric flux *through a flat surface placed perpendicular to a uniform electric field is the product of the field magnitude and the surface area.* It is a **scalar quantity**, and *it measures the number of field lines that cross the surface.* The SI unit of flux follows directly from the formula: newton metre-squared per coulomb, $\mathrm{Nm^2/C}$.

Now tilt the surface so that its normal makes an angle θ with the field. The same number of field lines still cross the surface, but if we project the tilted surface onto a plane perpendicular to the field, we get a smaller projected area equal to $A\cos\theta$. The flux therefore picks up a cosine factor:

$$\boldsymbol{\Phi = EA\cos\theta}$$

The figure shows the geometry: the flat surface of area A, the unit normal n̂ at the centre of the surface, and the angle θ between the field and the normal. When $\theta = 0°$ the surface is perpendicular to the field and the flux is at its maximum value EA. When $\theta = 90°$ the surface is parallel to the field, no field lines cross it, and the flux is zero.

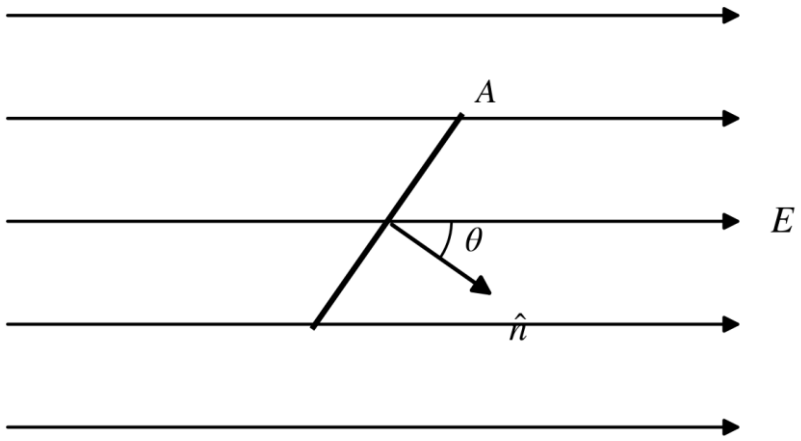


Figure: *Electric flux through a flat surface of area A tilted so that its unit normal n̂ makes angle θ with the uniform electric field E.*

A real surface need not be flat, and a real field need not be uniform. To handle the general case we break the surface into many tiny patches, each so small that the surface is essentially flat across the patch and the field is essentially uniform across it. Let dA denote the area of one such patch, and let its outward normal make angle θ with the local field. The patch contributes a small flux $d\Phi = E\,dA\cos\theta$. Summing the contributions of all the patches is, in the limit of infinitely small patches, a surface integral:

$$\Phi = \int E\,dA\cos\theta$$

If the surface is closed (an imaginary balloon with no hole in it), the integral is written with a small circle through the integral sign to remind us of the closure, and the convention is that the normal always points outward:

$$\Phi = \oint E\,dA\cos\theta$$

The electric flux through a surface *is the surface integral of the component of the electric field along the outward normal.* **For a closed surface**, *the integral counts the net number of field lines leaving the surface.* Lines that enter and then leave the surface contribute zero to the total, because their entry counts negatively and their exit counts positively, and the two cancel. This last property is the heart of the law we are about to state.

A word about field lines and space: A point charge sends its field lines out into space isotropically, meaning equally in every direction. The total angle of space around any point is 4π steradians (the three-dimensional analogue of the 2π radians around a point in a plane). The field lines therefore spread their entire share through that total solid angle. Now imagine wrapping any closed surface, of any shape, around the point charge. Every line that leaves the charge must, before it escapes to infinity, pierce that surface exactly once on its way out. The shape of the surface does not matter, only whether the charge is enclosed by it or not. This piece of pure geometry is what makes the law we are about to state independent of the shape of the surface we choose.

Gauss's law

The German mathematician Carl Friedrich Gauss, working in the early 1800s on problems in geodesy and electromagnetism, proved a beautiful result that is now one of the four equations of classical electromagnetism. In the language of electric flux it reads:

$$\oint \mathbf{E}\,\mathbf{dA}\cos\boldsymbol{\theta} = \frac{\mathbf{Q_{enc}}}{\boldsymbol{\varepsilon_0}}$$

In words, **Gauss's law** states that: *The net electric flux through any closed surface is equal to the total charge enclosed by the surface divided by the permittivity of free space.* The closed surface is called a **Gaussian surface**. It is imaginary; you draw it in space wherever it is convenient. You may choose any shape that makes the integral easy to evaluate, and the right-hand side will still equal $\frac{Q_{enc}}{\varepsilon_0}$ no matter what shape you pick, because the shape contributes only its enclosed charge to the answer.

Kipanga: *A surface I get to choose. Anywhere. Any shape. That feels too good to be true.*

Mr. Akilikubwa: *It is true, and like every gift in physics it comes with one condition. The right-hand side is always* $\frac{Q_{enc}}{\varepsilon_0}$*, regardless of shape; but the left-hand side, the integral, is only easy to evaluate when the shape you choose matches the symmetry of the charge distribution. For a single point charge you should choose a sphere centred on the charge. For an infinite line of charge you should choose a cylinder coaxial with the line. For an infinite plane you should choose a small pill-box piercing the plane. The art of using Gauss's law is choosing the right Gaussian surface.*

From Coulomb's law to Gauss's law

Gauss's law is not an extra postulate; it is a consequence of Coulomb's law and the geometry of three-dimensional space. To see this, place a single point charge Q at the centre of an imaginary sphere of radius r, and compute the flux of its electric field through that sphere. The geometry is shown in the figure.

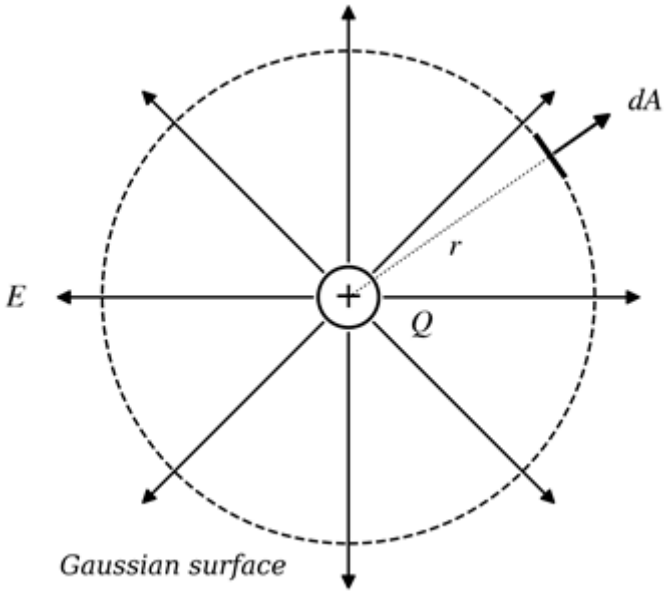


Figure: ***A Gaussian sphere of radius r around a point charge Q.*** *The radial field E is parallel to the outward area vector dA at every patch of the sphere, so cos θ = 1 everywhere on the surface.*

By the symmetry of the situation, the electric field on the sphere has the same magnitude everywhere on the surface, and it points radially outward at every point. Coulomb's law gives its magnitude:

$$E = \frac{kQ}{r^2}$$

The outward normal dA at every patch of the sphere is parallel to the local field, so the angle θ between them is zero everywhere, and $\cos\theta = 1$. The flux integral therefore simplifies to a constant E multiplied by the surface area of the sphere:

$$\Phi = \oint E\, dA \cos\theta = E \oint dA = E \times 4\pi r^2$$

Substitute the Coulomb expression for E:

$$\Phi = \frac{kQ}{r^2} \times 4\pi r^2 = 4\pi kQ$$

The radius r has dropped out. Now use $k = \frac{1}{4\pi\varepsilon_0}$ to replace k:

$$\boldsymbol{\Phi = 4\pi \times \frac{Q}{4\pi\varepsilon_0} = \frac{Q}{\varepsilon_0}}$$

The factor 4π inside the permittivity exactly cancels the 4π from the sphere's surface area, and we are left with the enclosed charge divided by ε_0, which is the right-hand side of Gauss's law. The geometric reason for the cancellation is the one already given in the solid-angle paragraph: the field lines of a point charge spread through 4π steradians, and ε_0 was historically defined with a 4π built into it so that the cancellation would happen cleanly.

From Gauss's law back to Coulomb's law

The argument runs equally well in reverse. Take Gauss's law as the starting point, place a single point charge Q alone in space, and surround it with an imaginary sphere of radius r centred on the charge. By symmetry the field on the sphere has the same magnitude E everywhere on the sphere, and points radially. The flux integral simplifies as before to $E \times 4\pi r^2$, and Gauss's law sets this equal to $\frac{Q}{\varepsilon_0}$:

$$E \times 4\pi r^2 = \frac{Q}{\varepsilon_0}$$

Make E the subject:

$$\boldsymbol{E = \frac{Q}{4\pi\varepsilon_0 r^2} = \frac{kQ}{r^2}},$$

which is Coulomb's law for the field of a point charge. The two laws are therefore equivalent: knowing one of them, you can derive the other. They look different because they encode the same physics from different starting points. Coulomb's law is a statement about pairs of charges and the force between them; Gauss's law is a statement about flux through closed surfaces. The next several pages will show why the flux version, when the geometry is symmetric, is the easier tool to reach for.

The flux recipe is laid out. One worked example arrives now to test the cosine factor before the four canonical applications of Gauss's law begin to fall, one by one, in the rest of the section.

BINDER Example 12

A flat circular surface of radius 20cm is placed inside a uniform electric field of magnitude 500N/C. The unit normal of the surface makes an angle of 60° with the field. Calculate the electric flux through the surface.

Solution

Let A be the area of the flat surface, E the magnitude of the uniform field, and θ the angle between the normal of the surface and the field. The geometry of the problem is shown in the figure.

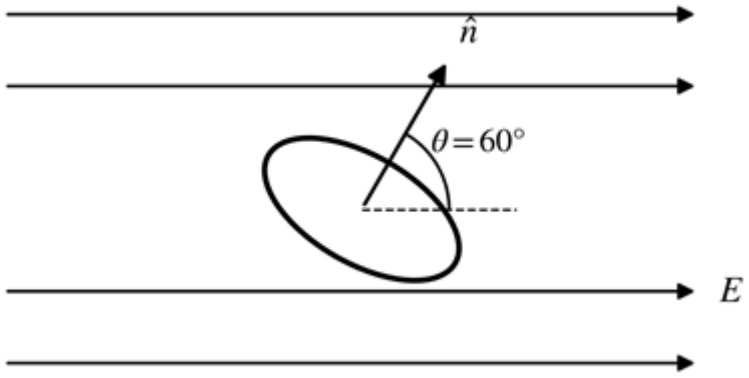


Figure: *A flat circular surface tilted so that its unit normal $\hat{n}$ makes an angle of 60° with the uniform electric field E.*

The area of a circular surface of radius r is $A = \pi r^2$. The electric flux through a flat surface in a uniform field is given by:

$$\Phi = EA\cos\theta = E\pi r^2\cos\theta$$

Substitute the data:

$$\Phi = (500\text{N/C}) \times \pi \times (0.20\text{m})^2 \times \cos 60° = 31.4\ \text{Nm}^2/\text{C}$$

The electric flux through the surface is $\Phi = 31.4\ \text{Nm}^2/\text{C}$

Making Sense of the Answer: *A quick sanity check. If the surface were perpendicular to the field ($\theta = 0°$), the flux would be $EA = (500N/C) \times \pi \times 0.040m^2 = 20\pi\ Nm^2/C = 62.8\ Nm^2/C$. If it were parallel to the field ($\theta = 90°$), no field lines would cross it and the flux would be zero. Our answer of $31.4\ Nm^2/C$ is exactly half of the maximum, which matches $\cos 60° = 0.5$ to within the precision of the arithmetic. The cosine is doing the geometric work for us.*

Think Like a Physicist: *Flux is a scalar that counts the net number of field lines crossing a surface, weighted by how perpendicular the surface is to those lines. The cosine factor is not a physical factor (nothing changes about the field or the surface), it is a geometric factor that accounts for the projection of the surface onto a plane perpendicular to the field. Once you have this idea in your hand, you will recognise it again in fluid flow, heat flow, and even in optics, because the same picture underlies every conservation law that talks about something flowing across a boundary.*

The flux recipe and Gauss's law are now in hand. Four canonical applications will occupy the next several pages. In each case the trick is the same: pick a Gaussian surface whose shape matches the symmetry of the charge distribution, observe that the field is constant on that surface, and read the magnitude of the field straight off the integral. The first application is the most fundamental one of all, because the charge distributions used in many later examples are built out of spherical pieces: a uniformly charged sphere.

1. A uniformly charged sphere

The uniformly charged sphere may be insulating or conducting.

A uniformly charged insulating sphere

Consider a solid insulator (for example a ball of plastic, or a hard rubber sphere) of radius R, carrying a total charge Q spread evenly throughout its volume. By **volume charge density** we mean the *charge per unit volume:*

$$\boldsymbol{\rho = \frac{Q}{\frac{4}{3}\pi R^3} = \frac{3Q}{4\pi R^3}}$$

For a uniformly charged sphere of radius R and total charge Q, the volume charge density is the same number ρ everywhere inside the sphere, and zero outside. To find the electric field at a point at distance r from the centre, draw an imaginary spherical Gaussian surface of radius r centred on the charged sphere. By the spherical symmetry of the problem, the magnitude of the field is the same everywhere on the Gaussian surface and points radially. The flux integral therefore simplifies as in the point-charge derivation:

$$\Phi = E \times 4\pi r^2 = \frac{Q_{enc}}{\varepsilon_0}$$

Two cases follow, depending on whether the Gaussian surface lies outside the charged sphere or inside it. The geometry is shown in the figure.

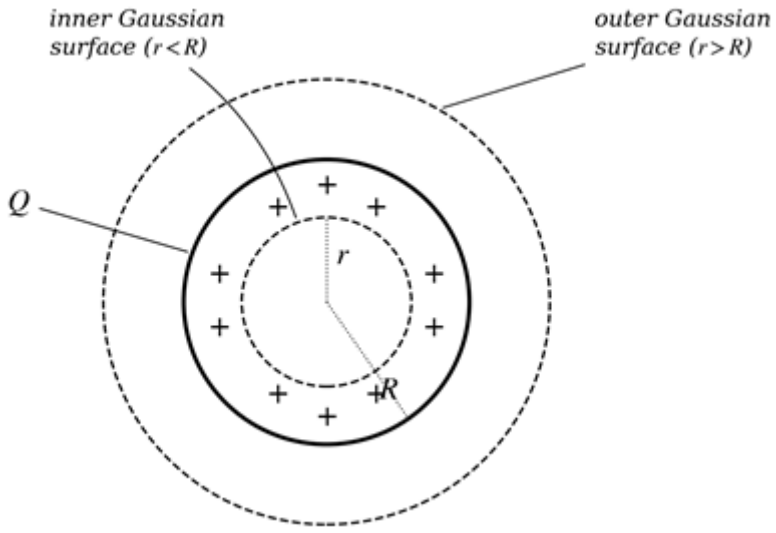


Figure: *A uniformly charged insulating sphere of radius R and total charge Q, with an inner Gaussian surface (r < R) and an outer Gaussian surface (r > R). The + signs indicate that the charge is spread throughout the volume of the sphere.*

Outside the sphere (r > R): When the Gaussian surface lies outside the charged sphere, it encloses the entire charge Q. Gauss's law gives:

$$E \times 4\pi r^2 = \frac{Q}{\varepsilon_0}$$

Make E the subject:

$$E = \frac{Q}{4\pi\varepsilon_0 r^2} = \frac{kQ}{r^2}$$

Outside the sphere, the field is exactly the same as if the entire charge Q were concentrated at the centre. This is the spherical analogue of Newton's shell theorem for gravitation, which we met in ***Foundations of Physics:*** a spherically symmetric shell or ball of mass attracts an external test mass as though all of its mass sat at its centre.

Inside the sphere (r < R): When the Gaussian surface lies inside the charged sphere, it encloses only the portion of the charge contained in the smaller sphere of radius r. That charge equals the volume charge density ρ multiplied by the volume of the inner sphere:

$$Q_{enc} = \rho \times \frac{4}{3}\pi r^3 = \frac{3Q}{4\pi R^3} \times \frac{4}{3}\pi r^3 = Q\frac{r^3}{R^3}$$

The factors of $4\pi/3$ have cancelled, leaving a clean cube ratio. Substitute this enclosed charge into Gauss's law:

$$E \times 4\pi r^2 = \frac{Q\frac{r^3}{R^3}}{\varepsilon_0}$$

Make E the subject and simplify:

$$\mathbf{E = \frac{Qr}{4\pi\varepsilon_0 R^3} = \frac{kQr}{R^3}}$$

Inside the sphere, the field grows **linearly** with the distance from the centre. It is zero at the centre ($r = 0$), and rises to its peak value $\frac{kQ}{R^2}$ at the surface ($r = R$). Outside the sphere, the field then falls as $\frac{1}{r^2}$. The same value at the surface is approached from both sides, so the field is continuous there. (We will see in a moment that the conducting case is different: it has a jump discontinuity at the surface.)

A uniformly charged conducting sphere

Replace the insulator with a metal sphere of the same radius R and the same total charge Q. A metal contains a sea of free electrons that can move across the bulk of the material. In electrostatic equilibrium (charges no longer moving), the electric field inside the metal must be zero, because if it were not zero the free electrons would feel a force and continue to drift, contradicting equilibrium. The verbal statement of this foundational result is short and worth noting cleanly:

Inside a conductor in electrostatic equilibrium, *the electric field is zero everywhere. As a consequence, any net charge placed on an isolated conductor resides entirely on its outer surface; there can be no net charge in the interior bulk.*

The figure shows the geometry: the same Gaussian surfaces as before, but now the charge Q lives on the rim of the sphere instead of being distributed through its volume.

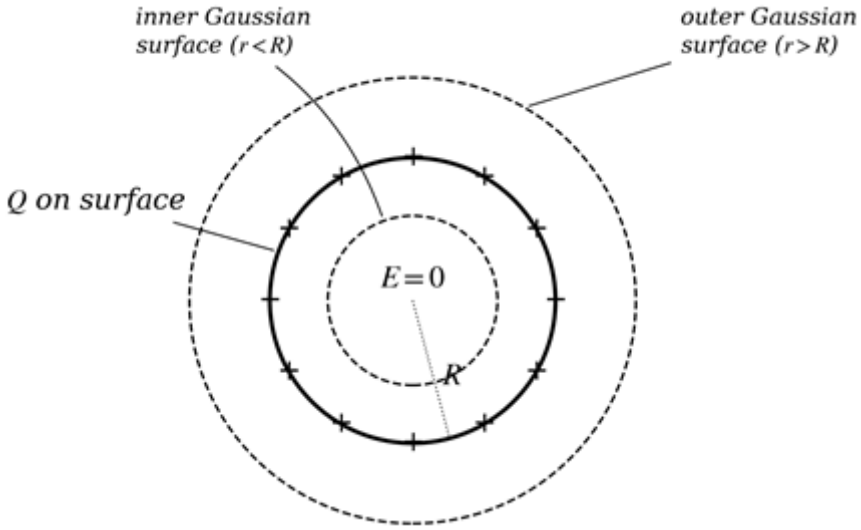


Figure: ***A uniformly charged conducting sphere of radius R with total charge Q.*** *The charge resides entirely on the outer surface (shown by the + signs along the rim), and the electric field is zero everywhere inside.*

Inside the conductor (r < R): The inner Gaussian surface encloses no charge at all, because the conductor's charge sits on the outer surface. Gauss's law therefore reads:

$$E \times 4\pi r^2 = \frac{0}{\varepsilon_0} = 0,$$

which gives $E = 0$ everywhere inside the conductor. This is the Gauss's-law derivation of the same result we just stated verbally: the interior of a conductor in electrostatic equilibrium is electric field-free.

Outside the conductor (r > R): The outer Gaussian surface encloses the full charge Q (all of which sits on the conductor's outer surface). Gauss's law reads exactly as for the insulating sphere:

$$E \times 4\pi r^2 = \frac{Q}{\varepsilon_0},$$

and the answer is the same:

$$E = \frac{kQ}{r^2}$$

Outside the conductor, the field is identical to that of either a point charge Q at the centre or the uniformly charged insulating sphere we studied a moment ago. From a distance, you cannot tell a metal sphere with charge Q from a plastic sphere with charge Q from a point charge Q: all three look the same to the field.

Insulating and conducting spheres, side by side

Outside the sphere, the two cases give the same answer. The difference lives inside. The figure plots the field magnitude as a function of distance r from the centre, for both spheres on the same axes.

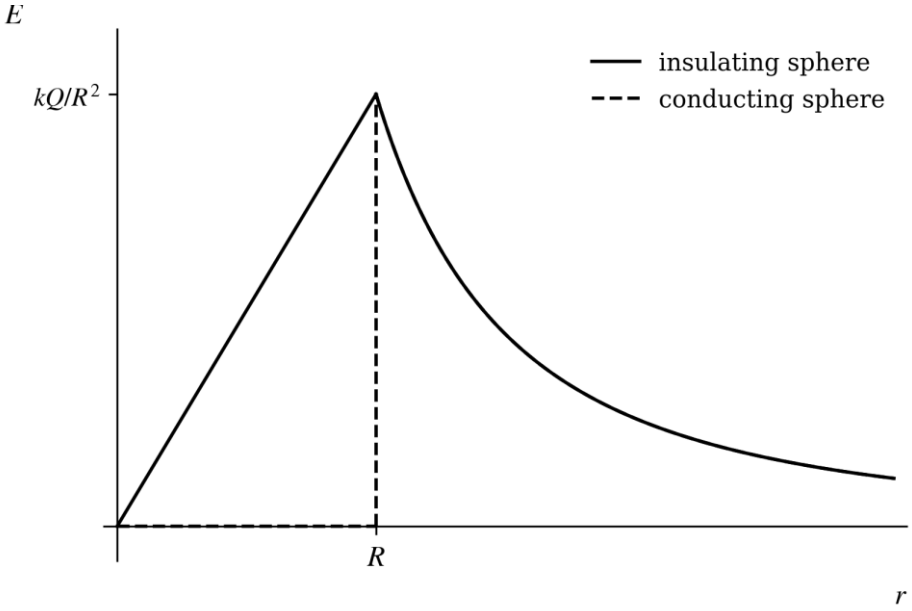


Figure: *Electric field magnitude E as a function of distance r from the centre, for a uniformly charged insulating sphere (solid line) and a conducting sphere of the same radius and charge (dashed line). Outside the sphere both curves coincide; inside, the insulating case grows linearly while the conducting case stays at zero.*

Two observations are worth committing to memory.

- The two curves are identical for $r > R$, so from outside the sphere we cannot detect by an electric measurement alone whether the body is an insulator or a conductor.
- The conducting case has a step discontinuity at the surface: the field jumps from zero just inside to $\frac{kQ}{R^2}$ just outside. The insulating case is smooth across the surface, because in that case the charge is spread through the volume rather than concentrated on a thin shell.

Kipute: *Sir, in the conducting case, where is the charge? You said it sits on the surface, but a surface has no thickness. So how much charge sits in a tiny patch of the surface?*

Mr. Akilikubwa: *A very good question. The natural quantity is the surface charge density, written with the Greek letter σ (sigma). For a conducting sphere of total charge Q and radius R, the charge is distributed evenly over the spherical surface of area $4\pi R^2$, so $\sigma = \frac{Q}{4\pi R^2}$. The field just outside the surface can then be written as $E = \frac{\sigma}{\varepsilon_0}$. We will derive this last result formally in the next application, when we look at an infinite plane sheet.*

The insulating-sphere result has two faces: the friendly $\frac{kQ}{r^2}$ on the outside, and the surprising linear-in-r face on the inside. The next example tests both faces in a single calculation, so that the contrast between them is visible in one place.

BINDER Example 13

An insulating sphere of radius 6.0cm carries a total charge of 3.0μC uniformly distributed throughout its volume. Calculate the magnitude of the electric field at a point 3.0cm from the centre, and at a point 9.0cm from the centre. Take the Coulomb constant to be $9.0 \times 10^9\ \mathrm{Nm^2/C^2}$.

Solution

Let Q be the total charge of the sphere, R its radius, and r the distance from the centre to the field point. The geometry is shown in the figure: one field point lies inside the sphere ($r < R$), the other lies outside ($r > R$).

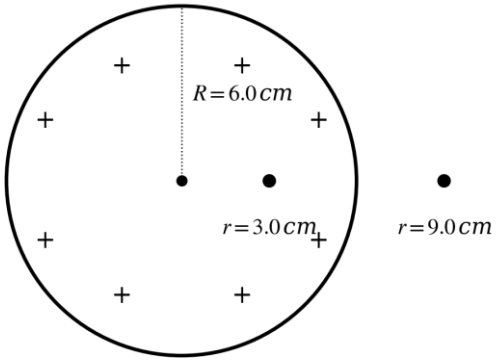


Figure: ***Example 13:*** *an insulating sphere of radius 6.0 cm carrying a total charge 3.0 μC. The field is wanted at the two indicated points.*

Field at the inside point: Here $r = 3.0\mathrm{cm} = 0.030\mathrm{m}$ is less than R, so the Gauss's-law result for the interior of a uniformly charged sphere applies:

$$E_{in} = \frac{kQr}{R^3}$$

Substitute the numerical values:

$$E_{in} = \frac{(9.0 \times 10^9\ \mathrm{Nm^2/C^2}) \times (3.0 \times 10^{-6}\ \mathrm{C}) \times (0.030\ \mathrm{m})}{(0.060\ \mathrm{m})^3} = 3.75 \times 10^6\ \mathrm{N/C}$$

Field at the outside point: Here $r = 9.0\mathrm{cm} = 0.090\mathrm{m}$ is greater than R, so the outside-the-sphere result applies, and the field is identical to that of a point charge Q at the centre:

$$E_{out} = \frac{kQ}{r^2}$$

Substitute the data:

$$E_{out} = \frac{(9.0 \times 10^9\ \mathrm{Nm^2/C^2}) \times (3.0 \times 10^{-6}\ \mathrm{C})}{(0.090\ \mathrm{m})^2} = 3.33 \times 10^6\ \mathrm{N/C}$$

The field at the inside point is $3.8 \times 10^6\ \mathrm{N/C}$, and the field at the outside point is $3.3 \times 10^6\ \mathrm{N/C}$.
The inside field is the larger of the two.

Making Sense of the Answer: *A moment of surprise is fair here. Naively, you might expect that getting closer to the centre always means a stronger field, because point-charge intuition suggests that the field should blow up as $1/r^2$ near the source. For a uniformly charged sphere, however, the field reaches its maximum at the surface, then falls off in both directions: linearly toward zero as you move toward the centre, and as $1/r^2$ as you move outward. The two evaluation points in this example sit at the same fractional distance on either side of the surface, $r/R = 0.5$ inside and $r/R = 1.5$ outside. The inside field is kQ/R^2 multiplied by $r/R = 0.5$, and the outside field is the same surface value multiplied by $(R/r)^2 = 1/2.25 \approx 0.44$. The inside number, 0.5, wins by a small margin over the outside number, 0.44, which is why the inside field came out slightly larger than the outside one.*

Think Like a Physicist: *The inside-versus-outside structure of the uniformly charged sphere is the spherical electrical twin of the Earth-as-ball gravitational problem from* ***Foundations of Physics:*** *inside a uniformly dense planet, the gravitational field grows linearly from zero at the centre to its peak value at the surface, then falls as $1/r^2$ outside. The two are not coincidences; they are two instances of the same theorem, which Newton proved for gravity in the seventeenth century and which Gauss proved in greater generality two centuries later. The inverse-square law and the geometry of three-dimensional*

space conspire to give a beautiful result that the field outside a spherically symmetric distribution depends only on the total amount enclosed, never on how that amount is arranged.

A uniformly charged hollow conducting sphere

The two spheres treated so far were solid all the way through. The case in which the metal is hollowed out, leaving an empty cavity inside an otherwise solid metal shell, is no less common, and far from being a curiosity it turns out to lie behind one of the most useful protective devices in all of electrical engineering. Take a hollow conducting sphere with inner radius R_1 and outer radius R_2, and let it carry a total positive charge Q deposited on it from outside, the whole arrangement left in electrostatic equilibrium.

The very first question is where the charge sits. The argument is the same as for the solid conductor and is worth repeating because it now does extra work. Inside the metal of any conductor in equilibrium the electric field must be zero, for if it were not, the free electrons of the metal would feel a force and the conductor would not yet be in equilibrium. Gauss's law applied to any closed surface lying wholly inside the metal then says that the charge enclosed must be zero too, since zero field gives zero flux. The total charge Q therefore cannot reside anywhere within the body of the metal; it can only sit on a surface. The metal of a hollow conductor has two surfaces, an inner one of radius R_1 facing the cavity and an outer one of radius R_2 facing the world. For a hollow conductor with no charge inside the cavity, all the charge migrates to the outer surface, and the inner surface carries none. This is the picture we work with.

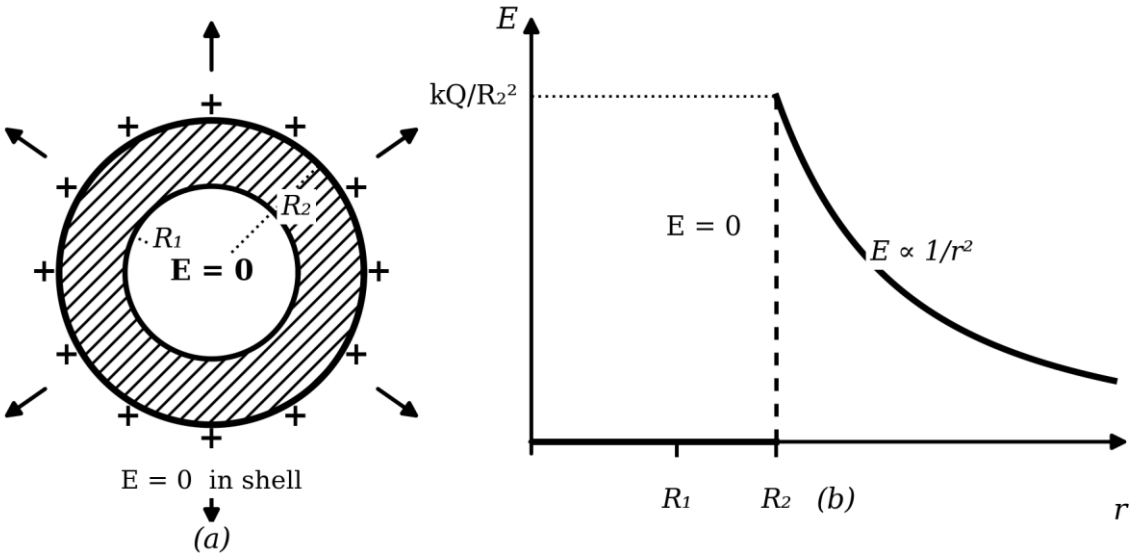


Figure: ***A uniformly charged hollow conducting sphere.*** *Panel (a): the conductor has inner radius R_1 and outer radius R_2. Charge sits entirely on the outer surface. The field is zero both in the cavity and within the metal of the shell. Outside the sphere the field is the same as if all the charge were concentrated at the centre. Panel (b): the field strength against distance. Zero everywhere inside, the field rises abruptly to its peak value kQ/R_2^2 at the outer surface, then falls as the inverse square of distance.*

With the charge sat down where it can be reached, Gauss's law is applied three times in turn, one Gaussian sphere for each of the three regions a radial coordinate can land in.

First region: Take first the region outside the conductor, $\mathbf{r > R_2}$. Draw a Gaussian sphere of radius r concentric with the conductor; it encloses the whole of the charge, Q. By the spherical symmetry of the problem the field is radial and has the same magnitude at every point of the Gaussian sphere, so the flux integral simplifies at once and Gauss's law reads:

$$E \times 4\pi r^2 = \frac{Q}{\varepsilon_0}$$

Solving for the field strength:

$$\mathbf{E = \frac{Q}{4\pi\varepsilon_0 r^2} = \frac{kQ}{r^2}}$$

Outside the hollow conductor the field is exactly the field of a point charge Q sitting at the centre. The hollowness makes no difference to the world beyond the outer surface, which sees only a total charge and the distance to it.

Second region: Take next the region within the metal of the shell itself, $\mathbf{R_1 < r < R_2}$. A Gaussian sphere of radius r drawn in this range lies entirely within the body of the metal, where the field is zero by the argument that opened the subsection. Zero field gives zero flux, and the calculation is complete in a single line: $\mathbf{E = 0}$ everywhere within the metal.

Third region: Take finally the region inside the cavity, $\mathbf{r < R_1}$. A Gaussian sphere of radius r drawn here encloses no charge whatever, because all of Q sits on the outer surface and lies outside this Gaussian. By Gauss's law the total flux through it is zero, and by the spherical symmetry the field on it is radial and of one fixed magnitude. The two facts can be reconciled only by setting that one magnitude to zero, and so the field inside the cavity is zero as well:

$$\mathbf{E = 0} \quad \text{for} \quad \mathbf{r < R_1}$$

Gathering the three answers together: the electric field is zero everywhere inside the hollow conductor, both within the metal of the shell and within the cavity it encloses. At the outer surface the field jumps abruptly to its peak value kQ/R_2^2, and from there outward it falls off as the inverse square of distance, exactly as if the whole charge were squeezed to a point at the centre. The two panels of the figure record this behaviour in picture and in graph.

The most striking thing about the result is the second region. A region of empty space, surrounded on all sides by a charged metal shell that may itself sit at a high potential, has zero field inside. The cavity is electrically calm regardless of what the shell is doing. This is the principle of the **Faraday cage**, and it is *why a person sitting inside a closed metal car during a lightning storm is safe even if the car itself is struck.* The metal body of the car becomes a charged hollow conductor for a brief instant, but the cavity inside it, where the occupants sit, holds its field-free peace. The very same result was leaned on in the previous section when the contrast between $E = 0$ and $V \neq 0$ was drawn, the cavity being the showcase example of a region where the field has vanished while the potential has not. Gauss's law has now done the work of placing that vanished field on a firm derivation.

A uniformly charged hollow insulating sphere

The fourth and final sphere of this section combines the hollow geometry of the third with the volume-filled charge distribution of the first. Take a hollow insulating sphere of inner radius R_1 and outer radius R_2, the shell material between the two radii carrying a uniform volume charge density ρ, the cavity inside the inner radius empty. Because the material is an insulator, the charge has no freedom to migrate; it stays where it was deposited and so really does fill the shell volume uniformly, instead of running to the outer surface as it did for the hollow conductor. The total charge of the sphere is then the volume charge density multiplied by the volume of the shell:

$$\mathbf{Q = \rho \times \frac{4}{3}\pi\left(R_2^3 - R_1^3\right)}$$

Again, Gauss's law is applied three times in turn, one Gaussian sphere for each of the three radial regions.

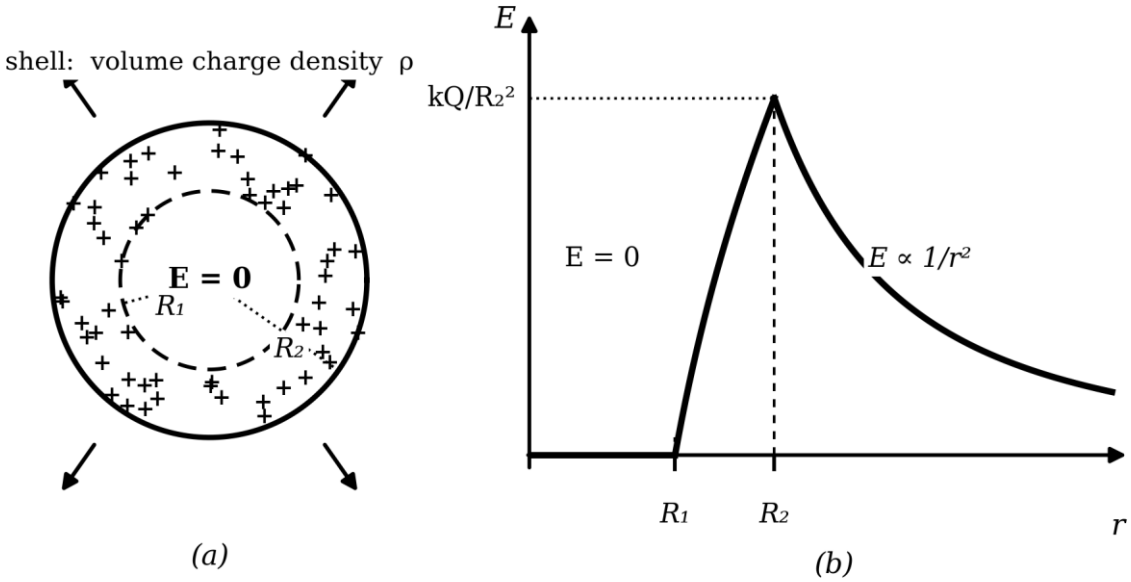


Figure: ***A uniformly charged hollow insulating sphere.*** *Panel (a): the insulating shell of inner radius R_1 and outer radius R_2 carries a uniform volume charge density ρ throughout its bulk. The cavity is empty. Panel (b): the field strength against distance. Zero inside the cavity, the field rises continuously through the shell from zero at the inner surface to its peak value kQ/R_2^2 at the outer surface, then falls as the inverse square of distance outside.*

First region: Take first the region inside the cavity, $\mathbf{r < R_1}$. A Gaussian sphere of radius r drawn here lies wholly inside the empty cavity and encloses no charge, since the charge is distributed entirely through the shell material lying outside this Gaussian. By Gauss's law the total flux is zero, and by the spherical symmetry the field on the Gaussian is radial and of one fixed magnitude. The two facts can only be reconciled by setting that magnitude to zero:

$$\mathbf{E = 0} \quad \text{for} \quad \mathbf{r < R_1}$$

The cavity is field-free for the same reason as in the hollow conductor of the previous treatment.

Second region: Take next the region within the shell of charge itself, $\mathbf{R_1 < r < R_2}$. A Gaussian sphere of radius r drawn in this range encloses only part of the charged shell, namely the portion of the shell that lies between R_1 and the radius r of the Gaussian itself. The volume of this enclosed portion is the volume of a spherical shell stretching from R_1 to r:

$$V_{enc} = \frac{4}{3}\pi (r^3 - R_1^3)$$

The charge inside the Gaussian is the density multiplied by this volume, and using the relation between Q and ρ above, the result can be written either as a density expression or as a simple fraction of Q:

$$\mathbf{Q_{enc} = \rho \times \frac{4}{3}\pi \left(r^3 - R_1^3\right) = Q \times \frac{\left(r^3 - R_1^3\right)}{\left(R_2^3 - R_1^3\right)}}$$

The ratio on the right is the fraction of the total charge that has been swept up by the Gaussian sphere of radius r, $\left(\frac{Q_{enc}}{Q} = \frac{r^3 - R_1^3}{R_2^3 - R_1^3}\right)$; the remaining charge sits in the still-unenclosed shell layer between r and R_2.

Gauss's law now reads:

$$E \times 4\pi r^2 = \frac{Q_{enc}}{\varepsilon_0}$$

Substituting Q_{enc} and solving for the field strength:

$$\mathbf{E = \frac{Q\left(r^3 - R_1^3\right)}{4\pi\varepsilon_0 r^2\left(R_2^3 - R_1^3\right)} = \frac{kQ\left(r^3 - R_1^3\right)}{r^2\left(R_2^3 - R_1^3\right)}}$$

This is the field inside the charged shell, and the three-cube structure of the result is the signature of this case. It does its checking at the two boundaries of the shell. At the inner boundary, $r = R_1$, the numerator factor $(r^3 - R_1^3)$ vanishes and the formula returns $E = 0$, matching the field already found in the cavity. At the outer boundary, $r = R_2$, the same factor becomes $(R_2^3 - R_1^3)$, which is exactly the factor sitting in the denominator and cancels with it, leaving $E = kQ/R_2^2$, the field of a point charge Q at the outer-surface distance. The shell formula has joined onto the cavity formula on its inside and joins onto the outside formula on its outside, the whole field varying continuously from one region to the next.

Third region: Take finally the region outside the sphere, $\mathbf{r > R_2}$. A Gaussian sphere of radius r drawn here encloses the whole of the charge, Q, and Gauss's law gives the by-now-familiar point-charge field:

$$\mathbf{E = \frac{kQ}{r^2}}$$

Gathering the three answers together: inside the cavity the field is zero, exactly as for the hollow conductor. Within the shell the field rises smoothly from zero at the inner surface R_1 to the peak value kQ/R_2^2 at the outer surface R_2, the rise governed by the $(r^3 - R_1^3) / r^2$ combination so that the field is small near R_1 and approaches its maximum near R_2. Outside the shell the field falls as the inverse square of distance, exactly as for a point charge. Panel (b) of the figure draws this whole behaviour as a single continuous curve.

With this fourth case, every spherical charge distribution that an examination is likely to set has now been worked through: the solid insulating sphere, the solid conducting sphere, the hollow conducting sphere, and the hollow insulating sphere.

The sphere is behind us. The next canonical surface is the one a sphere can never become, however large you make it: perfectly, endlessly flat. It is also the surface that produces the single most useful field in all of electrostatics, so it is worth meeting slowly.

2. An infinite plane sheet of charge

Imagine an enormous flat sheet, far wider than anything you intend to do near it: think of a great charged wall, or one face of the metal plates inside a capacitor, with charge spread evenly across it. We measure how thickly the charge is spread with the **surface charge density**, written **σ**, defined as the *charge carried on each unit of area of the sheet*. Its unit is the coulomb per square metre, C/m^2. A sheet with a large σ is a crowded sheet; a sheet with a small σ is a sparse one.

Before writing a single equation, stand at a point a short distance in front of this sheet and ask what the field there can possibly look like. The sheet is so vast that, from where you stand, it has no edges and no centre, and it looks exactly the same in every direction you might slide along it. A field cannot do something that the situation gives it no reason to do, and two firm conclusions follow from that thought alone.

First, the field must point straight out of the sheet, at right angles to it. Suppose for a moment that it leaned a little to one side. That sideways lean would have to choose one particular direction along the sheet, and yet every direction along the sheet is identical, so nothing could make the field prefer one over another. The only honest sideways component is no sideways component at all. The field is purely perpendicular to the sheet.

Second, the field has the same strength at every point that lies the same distance from the sheet, and it points directly away from the sheet on both faces if the sheet is positively charged, and directly toward it on both faces if the sheet is negative. A sheet with no special place on it cannot produce a strong field at one spot and a weak field at another the same distance away.

That is the picture. Now we choose a Gaussian surface that respects it, and Gauss's law hands us the magnitude of the field almost for free.

The natural surface to choose is a short, *closed cylinder*, often called a **pillbox**, pushed through the sheet so that it straddles it. Its two flat circular ends, each of area A, sit parallel to the sheet, one end poking out of each face, while its curved side wall runs perpendicular to the sheet. The figure shows the arrangement.

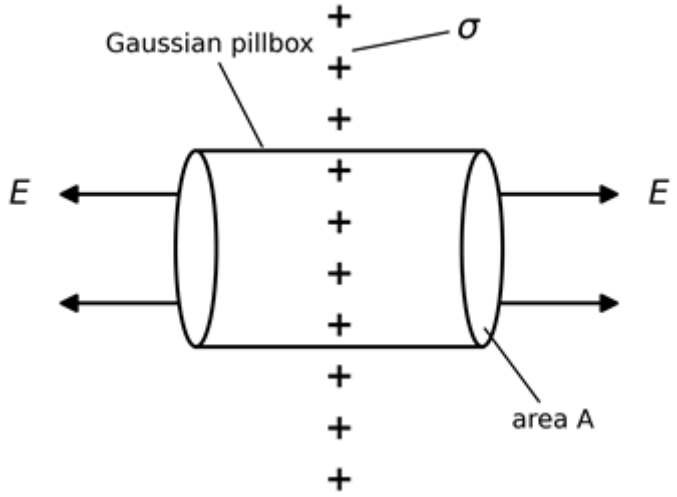


Figure: *A short closed cylinder, or pillbox, pushed through an infinite plane sheet of charge so that it straddles the sheet. The two flat caps, each of area A, lie parallel to the sheet, one on each side; the curved wall is perpendicular to it. The electric field E leaves both caps at right angles to the sheet, and the pillbox encloses a patch of the sheet of area A.*

Now we add up the electric flux passing out through this closed surface, piece by piece. Begin with the curved side wall. The field there points straight out of the sheet, while the wall itself runs perpendicular to the sheet, so the field merely skims along the wall without ever poking through it. The angle between the field and the wall's outward normal is a full right angle, and $\cos 90° = 0$, so the curved wall contributes no flux whatsoever.

That leaves the two flat end caps. On each cap the field is perpendicular to the surface and points straight out through it, exactly along the cap's outward normal, so $\cos 0° = 1$ and the flux through one cap is simply the field strength multiplied by the cap's area, EA. There are two caps, one on each side, and each contributes the same amount, so the total flux out of the whole pillbox is:

$$\Phi = \mathrm{EA} + \mathrm{EA} = 2\mathrm{EA}$$

Next we need the charge sealed inside the pillbox. The cylinder punches a small disc out of the sheet, a patch of area A, and whatever charge sits on that patch is the charge enclosed. Since σ is the charge on each unit of area, the charge on a patch of area A is $\mathrm{Q_{enc}} = \sigma\mathrm{A}$. No other charge lies inside the pillbox; the rest of the sheet is outside it.

Gauss's law states that the total flux out of a closed surface equals the charge it encloses divided by ε_0. Setting our two results equal to each other:

$$2\mathrm{EA} = \frac{\sigma \mathrm{A}}{\varepsilon_0}$$

Look closely at what happens next. The area A appears on both sides of the equation and cancels completely. Making E the subject, the field of an infinite plane sheet of charge stands revealed:

$$\mathbf{E} = \frac{\boldsymbol{\sigma}}{\mathbf{2\varepsilon_0}}$$

Since $\frac{1}{\varepsilon_0} = 4\pi\mathrm{k}$, the result may equally be written as:

$$\mathbf{E = 2\pi k\sigma},$$

whichever constant is the more convenient for the numbers in front of you.

Two features of this small formula deserve far more attention than its size suggests.

The first is that the area A vanished along the way. This is a relief rather than a surprise. The pillbox was an imaginary surface that we invented purely for the calculation, and a real, physical field cannot possibly depend on the size of an imaginary box. Had A survived into the answer, that would have been a sure sign of a mistake.

The second feature is genuinely startling. Search the formula $\mathrm{E} = \frac{\sigma}{2\varepsilon_0}$ for the distance from the sheet. It is not there. There is no r, no symbol of any kind for how far the chosen point lies from the sheet. Step one centimetre away from an infinite charged sheet, or step one kilometre away, and provided the sheet still looks infinite from where you stand, the field is exactly the same strength. The field of an infinite sheet does not fade with distance at all.

It is worth seeing why, because the reason is pure geometry. The field of a point charge weakens as $\frac{1}{r^2}$ because its field lines spray outward in every direction and must spread themselves ever more thinly over the surface of an ever-larger sphere. The field of a long straight line of charge weakens more gently, as $\frac{1}{r}$, because its lines spread only over the curved surface of an ever-larger cylinder. An infinite sheet, however, sends its field lines straight out, perpendicular and perfectly parallel to one another, and parallel lines never spread. With nothing to dilute them, the field they represent never weakens. A point charge spreads its lines over a sphere, a line spreads them over a cylinder, and a plane does not spread them at all.

What we have built, then, is a perfectly **uniform field**: the same strength and the same direction at every point on a given side of the sheet. Notice also the 2 sitting in the denominator. The field is $E = \frac{\sigma}{2\varepsilon_0}$ and not $\frac{\sigma}{\varepsilon_0}$ because the sheet fires its field out of both of its faces at once. The charge on the enclosed patch must share its flux between the two caps, and each side receives exactly half. That modest 2 is worth holding on to: the very next surface we build will remove it, and a great deal of useful physics lives in that difference.

A formula with no distance in it has a faintly suspicious look about it, rather like a train timetable printed with no times. One honest set of numbers will now be fed into it, and the suspicion will turn out to be the entire lesson.

BINDER Example 14

A large flat sheet carries a uniform surface charge density of $4.0\mu C/m^2$. Find the electric field at a point $2.0cm$ from the sheet, and then state the field at a point $50cm$ from the sheet. Take $\varepsilon_0 = 8.854 \times 10^{-12} C^2/(Nm^2)$.

Solution

The field of an infinite sheet of charge does not depend on the distance from the sheet, so a single calculation will answer both parts of the question at once. The field of a uniformly charged sheet is:

$$E = \frac{\sigma}{2\varepsilon_0}$$

Express every quantity in SI units. The surface charge density is $\sigma = 4.0\mu C/m^2 = 4.0 \times 10^{-6} C/m^2$. Substituting this value, together with ε_0:

$$E = \frac{4.0 \times 10^{-6} C/m^2}{2 \times 8.854 \times 10^{-12} C^2/(N \cdot m^2)} = 2.3 \times 10^5 N/C$$

The field at the point $2.0cm$ from the sheet is therefore $2.3 \times 10^5 N/C$.

The field at the point $50cm$ from the sheet is the very same value, $2.3 \times 10^5 N/C$, because the formula holds no distance for the change in position to act upon.

Making Sense of the Answer: *Two different distances were quoted in the question and a single field came back for both of them, which is not an oversight but the whole point of the example. It is worth pausing to feel how strange that is. Had the charge been a point charge instead of a sheet, moving from $2.0cm$ out to $50cm$, a twenty-five-fold increase in distance, would have weakened the field by a factor of twenty-five squared, that is by 625, dimming it almost to nothing. The sheet does no such thing. Its field lines leave the surface perpendicular and travel side by side without ever spreading, so there is nothing to thin the field out, and the needle of a field meter would not so much as twitch as you walked away. One honest caution belongs here for the real world: a genuine sheet has edges, and the result $E = \frac{\sigma}{2\varepsilon_0}$ is trustworthy only while the distance from the sheet stays small compared with the width of the sheet itself.*

Think Like a Physicist: *When a formula for a field has no distance in it, treat that empty space as a message rather than as a missing term. The absence is telling you that the geometry has arranged the field lines so that they never spread. This is a habit worth building: read a formula for what it leaves out just as carefully as for what it puts in. Here the three great charge shapes line up into one clean lesson. A point charge spreads its lines over a sphere, and its field falls as $\frac{1}{r^2}$. A long straight line spreads its lines over a cylinder, and its field falls as $\frac{1}{r}$. An infinite plane sends its lines out flat and parallel, they spread over nothing, and its field falls off not at all. Sphere, cylinder, plane: each shape writes its own signature into the way the field fades, and a physicist learns to read that signature straight off the formula.*

A field that flatly refuses to weaken with distance is an odd creature to be holding at the end of a worked example. Odd, but about to become extraordinarily useful. Keep this stubborn, uniform field firmly in mind, because the next surface we build does something clever with it: it brings in a second sheet, sets the two facing each other, and traps the field neatly in the gap between them.

3. Two parallel sheets of charge

Imagine two large flat sheets, each one an infinite sheet of the kind studied in the previous section, held parallel to each other with a narrow gap between them, like the two facing walls of a thin corridor. One sheet carries a positive surface charge density, which we write as $+\sigma$; the other carries an equal and opposite negative density, $-\sigma$. This is not an artificial

arrangement dreamt up to make a sum come out neatly. It is exactly the layout of charge inside a capacitor, the device on which every camera flash and every radio quietly depends, and which the next chapter will study in full.

Before any algebra, guess what the field of this pair ought to look like. The positive sheet pushes its field away from itself; the negative sheet pulls field toward itself. In the gap between the two, both of these urges point the same way, from the positive sheet across to the negative one, so we should expect a strong field there. In the regions outside the pair, the push of one sheet and the pull of the other are set against each other, so we should expect them to fight to a standstill and leave little or nothing behind. A strong, orderly field sealed inside the gap, and near-emptiness outside: that is the guess. Now we make it exact.

The pleasant surprise is that no new Gaussian surface is needed; the hard work was done already. A single infinite sheet, we found, produces a field of magnitude $\frac{\sigma}{2\varepsilon_0}$ on each side of itself, pointing directly away from the sheet if it is positive and directly toward the sheet if it is negative. Our arrangement is simply two such sheets, and the field at any point is the vector sum of what each sheet produces there on its own. Adding the separate fields of separate sources is called the **principle of superposition**, and it turns a fresh problem into a short piece of arithmetic.

Set the positive sheet on the left and the negative sheet on the right, and agree to call the direction running from the positive sheet toward the negative one, that is, rightward, the positive direction. Each sheet on its own contributes a field of size $\frac{\sigma}{2\varepsilon_0}$; all we must do is work out, in each region, which way those two contributions point, and then add them. The figure shows the arrangement and the field it produces.

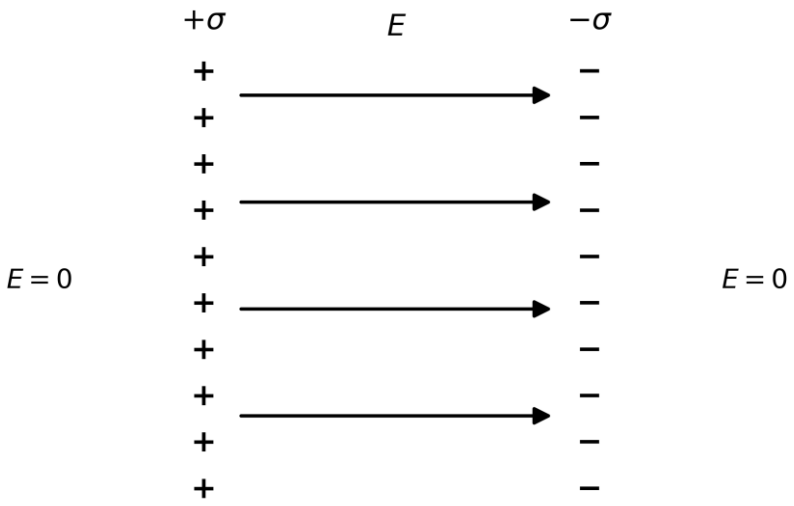


Figure: ***Two parallel sheets carrying equal and opposite surface charge densities, $+\sigma$ and $-\sigma$.*** *In the gap between the sheets the two single-sheet fields point the same way and reinforce, giving a uniform field E directed from the positive sheet to the negative one. In the regions outside the pair the two fields oppose and cancel, leaving $E = 0$.*

Begin **in the gap between the plates.** Here the field of the positive sheet points away from it, which is to the right. The field of the negative sheet points toward it, which is also to the right. The two contributions point the same way, so they add:

$$\mathrm{E} = \frac{\sigma}{2\varepsilon_0} + \frac{\sigma}{2\varepsilon_0} = \frac{\sigma}{\varepsilon_0}$$

The two equal halves combine into a single whole. The electric field in the gap between two parallel, oppositely charged sheets is therefore:

$$\mathbf{E} = \frac{\boldsymbol{\sigma}}{\boldsymbol{\varepsilon_0}}$$

Now step outside the pair, to the **left of the positive sheet**. The field of the positive sheet still points away from it, which out here means to the left. The field of the negative sheet still points toward it, which out here means to the right. The two contributions are equal in size but now point in opposite directions, so they cancel exactly:

$$\mathrm{E} = \frac{\sigma}{2\varepsilon_0} - \frac{\sigma}{2\varepsilon_0} = 0$$

The same cancellation occurs in the region to the right of the negative sheet, by an argument word for word the same. The electric field everywhere outside the pair of sheets is zero. The whole field of the arrangement has been swept into the gap and sealed there, exactly as the opening guess suggested.

It is worth pausing on the 2 that has just disappeared. In the previous section a single sheet gave a field $\frac{\sigma}{2\varepsilon_0}$, and the 2 in its denominator was there because that lone sheet had to share its flux between its two faces. Here, between the plates, that sharing works in our favour. Each sheet sends one half-strength field, $\frac{\sigma}{2\varepsilon_0}$, into the gap, and because the two half-fields point the same way they combine into a whole, $\frac{\sigma}{\varepsilon_0}$. The field between the plates is exactly twice the field of either sheet alone.

Notice what kind of field this is. Each single-sheet field was uniform, the same at every point on its side, and the sum of two uniform fields is uniform as well. The field in the gap has the same strength $\frac{\sigma}{\varepsilon_0}$ and the same direction at every point between the plates. It does not sag in the middle nor strengthen near a plate. This steady, even field is the one assumed, usually without a word of comment, in every parallel-plate problem a student meets, and it is the reason a charge released between the plates moves with a constant acceleration.

Notice too where the field is, and where it is not. It lies entirely inside the gap and is exactly zero everywhere outside. The pair of sheets has gathered its whole field into the narrow space between them and sealed it there. A capacitor is little more than this arrangement fitted with a pair of leads, and the field we have just found, $\mathrm{E} = \frac{\sigma}{\varepsilon_0}$, is the field inside every capacitor ever built. The next chapter takes this single result and, in no more than a few lines, turns it into the rule for how much charge such a device can store, $\mathrm{C} = \frac{\varepsilon_0 \mathrm{A}}{\mathrm{d}}$. A great deal of what that chapter does rests on the short derivation just completed.

The companion case: two sheets carrying the same sign

The arrangement just settled was the capacitor: two sheets carrying equal and opposite charge. It is worth to note what happens if the two sheets carry charge of the **same** sign instead, say $+\sigma$ on each, since the result is the exact mirror image of the capacitor's and explains in one stroke a piece of physics already met in this chapter.

Superposition does the whole calculation in three lines. Each sheet on its own contributes a field of magnitude $\frac{\sigma}{2\varepsilon_0}$ on each of its sides, directed away from the positive sheet. Setting the two sheets side by side with the rightward direction as positive, the contributions in the three regions add as follows:

$$\textbf{Left of both}: \quad \mathrm{E} = -\frac{\sigma}{2\varepsilon_0} + -\frac{\sigma}{2\varepsilon_0} = -\frac{\sigma}{\varepsilon_0}$$

$$\textbf{Between}: \quad \mathrm{E} = +\frac{\sigma}{2\varepsilon_0} + -\frac{\sigma}{2\varepsilon_0} = 0$$

$$\textbf{Right of both}: \quad \mathrm{E} = +\frac{\sigma}{2\varepsilon_0} + +\frac{\sigma}{2\varepsilon_0} = +\frac{\sigma}{\varepsilon_0}$$

Between the two like-charged sheets the two single-sheet fields point opposite ways and cancel exactly; outside the pair they point the same way and reinforce. The picture is the mirror image of the capacitor: the field is now expelled from the gap and dumped to either side, and the gap itself sits at zero. Same-sign sheets seal **nothing** between them.

This little result is the planar twin of the hollow-conductor result met under Gauss's law. There, charges spread uniformly over the outer surface of a hollow metal sphere produced zero field everywhere inside the cavity; here, charges of the same sign spread over two parallel sheets produce zero field everywhere between them. The two cases are the same physics in different geometries: when charges of one sign are arranged symmetrically around a region, their fields meet in the middle pointing opposite ways and cancel. This is the deep reason the Faraday cage works, and the present arrangement is the simplest possible setting in which to see it.

A single sheet handed us a field that ignored distance altogether. Two sheets, set face to face, strike an even better bargain: a strong and tidy field in the gap, and nothing whatsoever outside. One honest set of numbers will show how generous the bargain is.

BINDER Example 15

Two large parallel plates are held a short distance apart, one carrying a uniform positive surface charge density and the other an equal and opposite negative density, each of magnitude $8.0\mu\mathrm{C/m^2}$. (a) Find the electric field in the gap between the plates. (b) State the electric field in the region outside the plates. Take $\varepsilon_0 = 8.854 \times 10^{-12}\mathrm{C^2/(Nm^2)}$.

Solution

(a) The field in the gap between two parallel, oppositely charged sheets is uniform and has the magnitude:

$$\mathrm{E} = \frac{\sigma}{\varepsilon_0}$$

Express the surface charge density in SI units: $\sigma = 8.0\mu\mathrm{C/m^2} = 8.0 \times 10^{-6}\mathrm{C/m^2}$. Substituting this value together with ε_0:

$$\mathrm{E} = \frac{8.0 \times 10^{-6}\mathrm{C/m^2}}{8.854 \times 10^{-12}\mathrm{C^2/(Nm^2)}} = 9.0 \times 10^5\mathrm{N/C}$$

The field in the gap is therefore $9.0 \times 10^5\mathrm{N/C}$, uniform throughout the gap and directed from the positive plate toward the negative plate.

(b) Outside the plates the field of the positive sheet and the field of the negative sheet are equal in size and opposite in direction, so they cancel. The electric field everywhere outside the pair of plates is **zero.**

Making Sense of the Answer: *It is worth comparing this gap field with the field of a single sheet carrying the same σ. One sheet on its own would give $\frac{\sigma}{2\varepsilon_0}$, which works out to $4.5 \times 10^5 N/C$; the pair gives $9.0 \times 10^5 N/C$, exactly twice as much. The second sheet has doubled the field in the gap, and it has done so without adding any charge at all to the point where the field is measured. The zero outside is every bit as important as the large value inside. A field meter carried right up to the back of either plate would read nothing, because the two sheets cancel each other there. The honest caution from the previous section still holds: real plates have edges, and the tidy results, $\frac{\sigma}{\varepsilon_0}$ inside and zero outside, hold well in the interior of a wide pair of plates but blur where the edges are.*

Think Like a Physicist: *The whole of this section was settled without a single new integral and without drawing one fresh Gaussian surface. We already owned the field of one sheet, so we simply placed two of them and added. When a new problem is plainly built out of pieces you have already solved, do not start again from the beginning; identify the pieces and add their effects. That habit is the principle of superposition, and it is one of the most powerful labour-saving devices in all of physics. Notice also the shape of the result: the two fields reinforce in one region and cancel in another, purely because of the directions in which they point. Fields adding where they agree and cancelling where they disagree is a pattern you will meet again and again across physics, and the arithmetic never grows harder than it was here.*

So the humble pair of sheets turns out to be quietly magnificent. You have just derived, in a handful of honest lines, the exact field inside every capacitor on Earth. The next chapter does something almost unfair with that result: it takes this one formula and, in three short lines, extracts from it the rule for how much electrical charge a capacitor can hoard. Keep $\mathrm{E} = \frac{\sigma}{\varepsilon_0}$ somewhere safe. One canonical surface still remains, and it is the thinnest of them all: not a sheet at all, but a single endless line of charge.

4. An infinite line of charge

Imagine an endlessly long, straight, thin line, a charged thread, or a very long charged wire seen from close up, carrying charge spread evenly along its whole length. We measure how thickly the charge is spread with the **linear charge density**, written **λ,** defined earlier as the *charge carried on each unit length of the line*. Its unit is the coulomb per metre, $\mathrm{C/m}$.

Stand at a point a perpendicular distance x from this line and ask, as we have asked twice before, what the field there is able to be. The line looks identical whether you face along it one way or the other, so the field can carry no component running along the line; any such component would have to choose a direction the line itself flatly refuses to distinguish. The line also looks identical from every direction around it. Two conclusions follow at once. First, the field points straight out from the line, radially, at right angles to it. Second, its strength depends only on x, the perpendicular distance from the line, and is the same at every point that shares that distance.

This is cylindrical symmetry, and it calls for a cylindrical Gaussian surface. We take an imaginary cylinder of radius x and length L, sharing its axis with the line of charge, so that the line runs straight down its centre. The figure shows the arrangement.

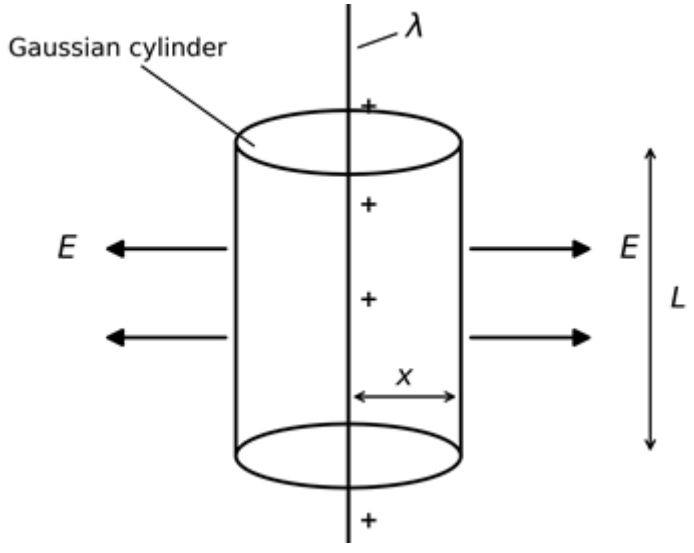


Figure: ***A cylindrical Gaussian surface of radius x and length L, sharing its axis with an infinite line of charge of linear density λ.*** *The field E points radially outward and pierces only the curved wall; it skims along the two flat end caps and so sends no flux through them.*

Add up the flux through this closed cylinder in two parts: its two flat end caps, and its curved side wall. Take the end caps first. On either cap the outward-pointing normal runs parallel to the line itself, while the field there points radially, square across the line. The field thus skims along the face of the cap without ever passing through it; the angle between the field and the cap's normal is a full right angle, $\cos 90° = 0$, and the two end caps together contribute no flux at all.

The curved side wall is another matter entirely. Everywhere on it the field points radially outward, and so does the wall's own outward normal; the two are exactly aligned, $\cos 0° = 1$. Better still, every point of the curved wall lies at the same perpendicular distance x from the line, so by our symmetry the field has the same strength E at every point of it. The flux through the curved wall is therefore E multiplied by the wall's area. A cylinder of radius x and length L has a curved wall of area $2\pi xL$, its circumference $2\pi x$ multiplied by its length L. The total flux out of the whole closed surface is then:

$$\Phi = E(2\pi xL)$$

The charge enclosed is whatever length of line lies sealed inside the cylinder. The cylinder is L long and the line runs straight down its axis, so it encloses a length L of charged line. With λ the charge on each unit of length, the enclosed charge is $Q_{enc} = \lambda L$.

Gauss's law now sets the total flux equal to the enclosed charge divided by ε_0:

$$E(2\pi xL) = \frac{\lambda L}{\varepsilon_0}$$

The length L stands on both sides of the equation and cancels, and well it should, for L was a feature of an imaginary cylinder, never of the real line. Making E the subject delivers the field of an infinite line of charge:

$$\mathbf{E = \frac{\lambda}{2\pi\varepsilon_0 x}}$$

Since $\frac{1}{2\pi\varepsilon_0} = 2k$, the same result is just as often written as:

$$\mathbf{E = \frac{2k\lambda}{x}},$$

whichever constant sits more comfortably with the numbers in front of you.

Look closely at how this field fades. It falls off as $\frac{1}{x}$: double your distance from the line and the field drops to one half. This sits neatly between the two fields found earlier in this section. The field of a point charge spreads its lines outward over the surface of a sphere, whose area grows as the square of the distance, so that field falls as $\frac{1}{r^2}$. The field of an infinite sheet sends its lines straight out, parallel, and never lets them spread, so that field does not fall away at all. A line lies exactly in between the two. Its lines spread over the curved wall of a cylinder, whose area grows only in proportion to x, not to its square, so the field falls as $\frac{1}{x}$, gently. Sphere, cylinder, plane; an inverse square, an inverse first power, and no falloff at all. The geometry of the source is written directly into the way its field dies away.

There is one more thing worth saying, and it is really the whole point of this section. Earlier in this chapter we found this very same field, $E = \frac{\lambda}{2\pi\varepsilon_0 x}$, by the long road. We chopped the line into countless tiny elements, wrote down the small field contributed by each, and summed them all with an integral. It worked, but it was heavy labour. Here, Gauss's law has reproduced the identical result in a single line of algebra, almost the moment the right cylinder was drawn. That is the power the law has been promising from the start. Choosing a Gaussian surface that matches the symmetry of the charge, a sphere for a point or a ball, a pillbox for a sheet, a cylinder for a line, turns a hard integral into one short multiplication. Picking the surface is the real physics; the algebra that follows is light work.

Four formulae for four shapes are now in hand. The line's formula is the only one of them that fades with distance in a gentle, well-mannered way, and a single calculation will show exactly how gentle that gentleness is.

BINDER Example 16

A long straight wire carries a uniform linear charge density of $5.0 \times 10^{-8}\text{C/m}$. (a) Find the electric field at a point 0.10m from the wire. (b) Hence state the field at a point 0.20m from the wire. Take $\varepsilon_0 = 8.854 \times 10^{-12}\text{C}^2/(\text{Nm}^2)$.

Solution

(a) The field of an infinite line of charge, at a perpendicular distance x from it, is:

$$E = \frac{\lambda}{2\pi\varepsilon_0 x}$$

Substitute λ, ε_0 and $x = 0.10\text{m}$, every quantity in SI units:

$$E = \frac{5.0 \times 10^{-8}\text{C/m}}{2\pi \times 8.854 \times 10^{-12}\text{C}^2/(\text{Nm}^2) \times 0.10\text{m}} = 9.0 \times 10^3\text{N/C}$$

The field at the point 0.10m from the wire is $9.0 \times 10^3\text{N/C}$, directed radially away from the wire.

(b) The field of a line of charge falls off as $\frac{1}{x}$, so doubling the distance halves the field. Without a fresh substitution, the field at 0.20m is simply half the value just found, $4.5 \times 10^3\text{N/C}$, again directed radially outward.

Making Sense of the Answer: *The two answers differ by exactly a factor of two for a doubling of distance, and that is the unmistakable signature of a* $\frac{1}{x}$ *field. It is worth holding the three signatures of this section side by side. Move twice as far from a point charge and its field drops to a quarter, since an inverse-square law falls away fast. Move twice as far from an infinite sheet and its field does not change in the slightest. Move twice as far from a line and the field drops to one half: the line is the moderate middle case of the three. Part (b) also makes a quiet point about method. Once the shape of a law is known, a second value can often be read straight off the first, with no need to repeat the whole calculation. The familiar real-world caution still belongs here: the formula treats the line as infinitely long, so it is trustworthy only for points whose distance from the wire is small compared with the wire's own length, comfortably clear of its ends.*

Think Like a Physicist: *This whole section asked one skill of you, over and over: match the imaginary surface to the symmetry of the real charge. A point or a sphere has spherical symmetry, so a spherical surface was the one to draw. A sheet has flat symmetry, so a flat-faced pillbox was drawn. A line has cylindrical symmetry, so a cylinder was drawn. In every single case, once the surface fitted the symmetry, the field stepped outside the flux as a plain constant and Gauss's law quietly did the rest. So when you meet a new charge distribution, do not reach for a formula first; reach for its symmetry. Ask what the field is forbidden to do, let that forbidding pin down both its direction and the surface that respects it, and the law will hand you the magnitude. Symmetry first, surface second, algebra last: that is the entire method in six words.*

And there it is: four shapes, four fields, one law. Gauss's law has now been aimed in turn at a sphere, a sheet, a facing pair of sheets and a line, and each time it has paid out a clean formula almost for the asking. One short calculation still waits to close the section, and it is a faintly dangerous one. It asks how much charge you may pile onto an object before the very air around it loses patience, breaks down, and answers you back with a spark.

That spark deserves a proper explanation, and, pleasingly, it calls for no new physics at all. The first of the four applications just completed, the charged sphere, joined to a single new fact about air, is enough to predict exactly when it happens.

Dielectric breakdown: when the air gives way

Air, in its ordinary state, is an insulator: it does not conduct electricity, which is why charge can sit undisturbed on objects and why the wires inside a wall do not simply discharge into the room. But this refusal to conduct has a limit. Drive a strong enough electric field through air and the field begins to tear electrons loose from the air molecules. Those freed electrons, hurled along by the same field, crash into further molecules and knock loose more electrons still, which are themselves hurled onward to free yet more. In a fraction of a second the air along that path has turned from an insulator into a conductor, and charge floods across it. We see and hear that flood as a spark. Lightning is the very same process, run on the scale of a thundercloud.

Every insulator has a ceiling: the strongest electric field it can withstand before it breaks down in this way and begins to conduct. That ceiling is called the **dielectric strength** of the material. For dry air at ordinary pressure the dielectric strength is about $3.0 \times 10^6 \mathrm{V/m}$, three million volts per metre. The volt per metre is simply another name for the unit of electric field; a field of one $\mathrm{V/m}$ is the same as a field of one $\mathrm{N/C}$. Keep the field in air below that figure and the air insulates in silence; let the field once reach it, and the air gives way.

Now the question that closes this section. Take a metal sphere of radius R carrying a charge Q, standing in open air. The very first of the four applications told us the field such a sphere produces: outside the sphere, the field is exactly that of a point charge Q placed at the centre. That field is strongest right at the sphere's own surface, a distance R from the centre, where its magnitude is $\frac{kQ}{R^2}$; further out it only weakens, falling away as $\frac{1}{r^2}$. As charge is added to the sphere, Q climbs and the surface field climbs with it. The sphere cannot be charged without limit. The moment its surface field reaches the dielectric strength of air, the air at the surface breaks down, and any further charge simply sprays off into the surroundings as **sparks.** The greatest charge the sphere can hold is therefore the charge that lifts its surface field exactly to that ceiling.

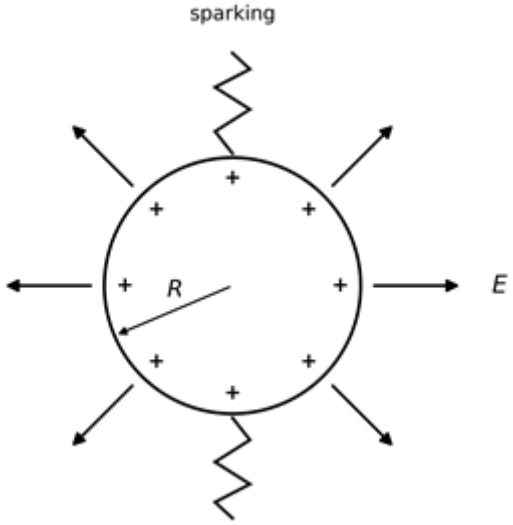


Figure: ***A charged conducting sphere of radius*** **R.** *The field* E *is strongest right at the surface and weakens outward. When the surface field reaches the dielectric strength of the surrounding air, the air breaks down and the sphere sparks; this sets a ceiling on the charge it can hold.*

Anyone who has tugged off a wool jumper on a dry day, and heard it crackle in the dark, has already run this experiment in miniature. Here it is again, this time with the numbers filled in.

BINDER Example 17

A metal sphere of radius 10cm stands in dry air, whose dielectric strength is $3.0 \times 10^6 \text{V/m}$. Find the greatest charge the sphere can hold before the air at its surface breaks down and it begins to spark. Take $\text{k} = 8.99 \times 10^9 \text{Nm}^2/\text{C}^2$.

Solution

The field just outside a charged sphere is the field of a point charge sitting at its centre. At the surface, a distance R from that centre, its magnitude is:

$$\text{E} = \frac{\text{kQ}}{\text{R}^2}$$

The air breaks down when this surface field reaches the dielectric strength. Setting the surface field equal to that ceiling and making Q the subject:

$$\text{Q} = \frac{\text{ER}^2}{\text{k}}$$

Substitute the dielectric strength for E, with $\text{R} = 10\text{cm} = 0.10\text{m}$, keeping every quantity in SI units:

$$\text{Q} = \frac{3.0 \times 10^6 \text{N/C} \times (0.10\text{m})^2}{8.99 \times 10^9 \text{N} \cdot \text{m}^2/\text{C}^2} = 3.3 \times 10^{-6}\text{C}$$

The greatest charge the sphere can hold is therefore about $3.3 \times 10^{-6}\text{C}$, a little over three microcoulombs.

Making Sense of the Answer: *Three microcoulombs is a startlingly small charge, a few millionths of a single coulomb, and yet it is the most that a sphere the size of a large grapefruit can hold in open air. The result is really a statement about the size of the coulomb itself: the coulomb is an enormous unit, and a whole coulomb of static charge resting on a small object in air is something nature simply will not permit. It is the reason the charges met all through electrostatics are almost always quoted in microcoulombs or nanocoulombs. The size of the sphere matters as well. Because the surface field is* $\frac{kQ}{R^2}$*, a larger sphere can take on more charge before that field reaches the ceiling, while a small sphere, or worse a sharp point with a very small radius of curvature, reaches the ceiling at a far smaller charge. That is precisely why high-voltage equipment is built with smooth, well-rounded surfaces, and why a sharp metal spike is the first thing to spark, which is the very principle a lightning conductor is built upon.*

Think Like a Physicist: *This problem arrived dressed as something new, the dramatic business of sparking air, and yet it was solved with nothing new at all. The whole calculation rested on a single result derived earlier in this section, the field outside a charged sphere, joined to one new fact about air. When a problem looks unfamiliar, the first move is to search it for the familiar piece hidden inside. Notice also the shape of the method. The question asked for a maximum, the most charge the sphere could hold, and the route to it was to take the formula for the relevant quantity, set it equal to its limiting value, and solve. A great many physical questions, the most, the least, the longest, the fastest, the moment just before something gives way, are answered by exactly that move: name the limit, set the formula equal to it, and solve.*

With that, Gauss's law has given, for the moment, all that it has to give. This section set out from a single idea, that counting the field lines crossing a closed surface counts the charge sealed within it, and from that one idea it has produced the field of

a sphere, of a sheet, of a facing pair of sheets and of a line, and finally the precise charge at which the surrounding air revolts. The electric field has now been mapped thoroughly. The next section turns to a different and equally powerful question: **no**t *how hard the field pushes at each point,* but *how much energy is stored in the pushing,* the idea of electric potential.

ELECTRIC POTENTIAL AND POTENTIAL DIFFERENCE

Until now the electric field has held the stage: a vector at every point of space, telling any charge placed there which way it will be pushed and how hard. That is one full and correct description of the electrical world. But there is a second description, and for a great many problems it is far the more convenient. It speaks not of forces and directions but of energy, and its central idea is **electric potential.**

To reach that idea, think first about what happens when a charge is moved. Carry a small positive charge from one point in a field to another. If you move it against the field, you must push, and in pushing you do work on it; that work does not vanish but is stored, ready to be released, exactly as the work done in lifting a stone is stored as gravitational potential energy. If instead you let the charge run along the field, the field itself does the work, and the charge speeds up, gaining kinetic energy as it goes. Either way, energy is bound up in a charge's position within a field. This is the thread we now follow.

Take two points in a field, label them A and B, and let an external agent carry a small positive test charge q slowly from B to A, slowly enough that it never picks up speed. Call the work the agent does on that journey W_{AB}. Now comes the key step: divide that work by the size of the charge moved. The quantity $\frac{W}{q}$ does not depend on how large or small a test charge was chosen, for doubling the charge doubles the work and leaves the ratio untouched. It is therefore a property of the two points and the field between them, and of nothing else. We give it a name, the **potential difference** between A and B and is expressed as the work done per unit charge in carrying a positive charge from B to A:

$$\mathbf{V_{AB} = \frac{W_{AB}}{q}}$$

In words: **The potential difference** *between two points is the work done per unit charge in carrying a positive charge from one of the points to the other.*

Because it is an energy divided by a charge, the unit of potential difference is the joule per coulomb. This unit is met so often that it is given its own name, the **volt**, after Alessandro Volta: one volt is one joule per coulomb, $1V = 1J/C$. To say that the potential difference between two points is one volt is to say that one joule of work is needed to carry each coulomb of charge between them, no more and no less.

That definition is the practical one, but it leaves the field itself out of sight. The two are easily joined. To carry the charge against the field, the agent must, at every step of the way, push with a force equal and opposite to the force qE $\left(\text{from } E = \frac{F}{q}, F = qE\right)$ that the field exerts on it. Adding up the agent's work, step by tiny step, along the path from B to A, gives the potential difference directly:

$$\mathbf{V_{AB}} = -\int_B^A \frac{qE \times dr}{q} = -\int_{\mathbf{B}}^{\mathbf{A}} \mathbf{E} \,.\, \mathbf{dr}$$

The integral simply sums the field's effect along the path; the dot is a reminder that only the part of the field lying along each small step does any work, and the minus sign records that moving a positive charge against the field calls for positive work from the agent. The field, put the other way round, points in the direction in which the potential falls.

One feature of this integral is too important to pass over. Its value depends only on the two endpoints, B and A, and not in the least on the route taken between them. A charge carried from B to A along a short straight path, and the same charge carried along a long and winding one, call for exactly the same work per unit charge. The figure shows two such routes. The reason is that the electrostatic force, like gravity, is a **conservative force**: the work it does around any closed loop is zero. This is what makes potential difference a sensible idea at all. Were the work to depend on the path, $\frac{W}{q}$ would not be a property of the two points, and could not be quoted as a single definite number.

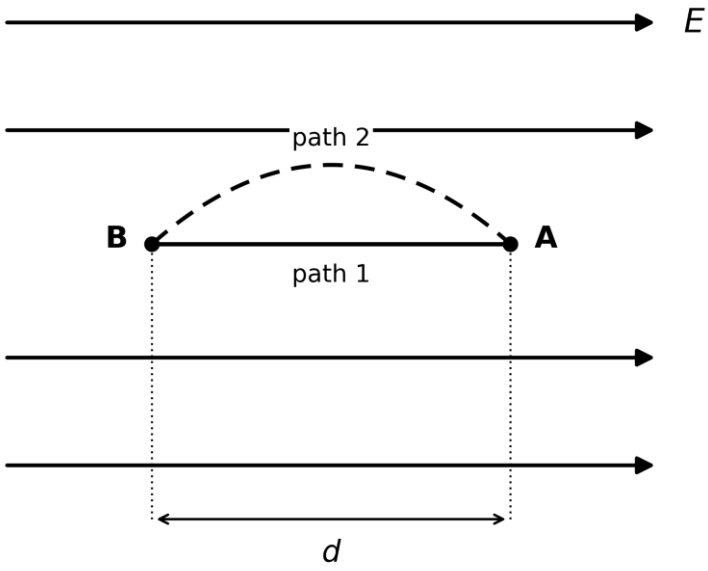


Figure: ***Two different routes, a straight one and a winding one, carrying a charge between the same two points*** **B** ***and*** **A** ***in a uniform field*** **E.** *The work done per unit charge is the same along both paths; it depends only on the endpoints, which is what allows a single potential difference to be assigned to the pair.*

Potential difference compares two points. It is often useful to speak instead of the potential at a single point, and for that we need one fixed point of reference to measure everything against, just as heights are quoted from sea level. The natural choice is a point infinitely far from all charges, where the field has faded to nothing; there the potential is taken to be zero. This results to the new term, **electric potential.**

The **electric potential** at a point is then *the work done per unit charge in bringing a positive charge from infinity to that point.*

The potential at a point near a positive charge is itself positive, since an agent must do positive work to push a positive test charge inward against the repulsion; the potential near a negative charge is negative. A potential difference is then simply the difference of two such potentials, $V_{AB} = V_A - V_B$.

And here is the great convenience promised at the outset. Electric potential is a **scalar**. At each point of space it is a single number, a count of joules per coulomb, with no direction attached to it. The field needed three numbers at every point, a strength and a direction; the potential needs only one. Problems that are awkward to handle with vectors often become straightforward once recast in terms of this single scalar, as the next section will show in full.

One case is worth working out at once, for it is the one met most often. Suppose the field is uniform, the same in strength and direction at every point, as it is in the gap between the charged parallel plates of the previous section. Carry a charge a distance **d** directly along such a field. The field strength E never changes along the way, so the sum that the integral performs is no harder than a single multiplication, and the potential difference between the two points is simply:

$$\mathbf{V = Ed}$$

The potential difference across a uniform field is the field strength multiplied by the distance moved along it. Travelling along the field, in the direction it points, the potential falls; travelling against the field, it rises. A uniform field of E volts per metre shifts the potential by E volts for every metre crossed along it. This single relation, $V = Ed$, together with the definition V_{AB} as work per unit charge, will carry us through most of what follows.

Definitions stay slippery until a number is poured into them. Here is a uniform field, a marked-off distance and a charge to be carried across it; two short questions will turn the definitions solid.

BINDER Example 18

A uniform electric field of magnitude $2.0 \times 10^3 \text{V/m}$ fills the region between two points A and B that lie a distance 5.0cm apart, measured along the field. (a) Find the potential difference between A and B. (b) Find the work that must be done to carry a charge of 3.0μC from one of the points to the other.

Solution

(a) The two points lie along a uniform field, so the potential difference between them is the field strength multiplied by the distance separating them, $V = Ed$.

Substituting the given values in SI units, with $d = 5.0\text{cm} = 0.050\text{m}$:

$$V = (2.0 \times 10^3 V/m)(0.050m) = 100V$$

The potential difference between A and B is $100V$.

(b) The work needed to carry a charge across a potential difference follows straight from the definition of that difference as work per unit charge, $W = qV$. With $q = 3.0 \times 10^{-6}C$ and the potential difference just found:

$$W = qV = (3.0 \times 10^{-6}C)(100V) = 3.0 \times 10^{-4}J$$

The work that must be done is $3.0 \times 10^{-4}J$, or $0.30mJ$.

Making Sense of the Answer: *Notice that part (b) never once needed the field, the distance, or the geometry of the situation. As soon as the potential difference between the two points was known, the work followed from* $W = qV$ *alone, by the very definition of potential difference as work per unit charge. Notice the size of the answer too. Carrying three microcoulombs through a hundred volts costs only three ten-thousandths of a joule, less than the work of lifting a grain of rice a finger's width. Everyday electrostatic energies are small, even when the voltages, hundreds or thousands of volts, sound large, simply because the charges in question are so tiny. One point of care: a potential difference has a sense as well as a size. The potential is the higher at the point reached by moving against the field. Carrying a positive charge from the low-potential point to the high one demands positive work from you; carry it the other way and the field does the work for you, handing the charge energy it can spend.*

Think Like a Physicist: *Two descriptions of electricity now sit side by side, and it is worth being clear about how they differ. The electric field is a vector: at every point it carries both a strength and a direction, and to combine fields means to add vectors, with all the care that arrows demand. The electric potential is a scalar: at every point it is one plain number. Both describe the same physics, and the relation between them, the line integral met above, or in a uniform field the tidy* $V = Ed$*, lets you pass freely from one to the other. The skill is to choose the easier language for the question in hand. If a problem asks which way a charge will move, the field, with its built-in direction, answers at once. If a problem asks how much energy a charge will gain, or how much work a process will cost, the potential, a single number needing no vector arithmetic, is almost always the quicker road. A physicist keeps both maps of the one territory, and reads whichever makes the present journey shorter.*

We have now defined the potential and learned to count it out in a uniform field. But uniform fields are the gentle exception. Most fields, the field of a point charge first among them, vary from place to place. The next task is to find the potential of a single point charge, and to discover, with it, a small miracle of convenience: where electric fields must be added as vectors, electric potentials, being mere numbers, may simply be added up.

The potential of a point charge

We ended the last section with a promise: that potentials, unlike fields, may be combined as plain numbers. To redeem it we need first the potential of the simplest source there is, a single point charge, and then the rule for adding several of them together.

Take a single point charge Q sitting alone in space, and ask for the potential it produces at a point a distance r away. That potential, as the last section established, is the work done per unit charge in carrying a small positive test charge to the point from infinitely far off, where the potential is counted as zero. The field of a point charge is already known to us: it points radially outward, with magnitude $E = \frac{kQ}{r^2}$. Because the test charge may be carried straight in along a radius, the field lies squarely along every step of the path, and the line integral of the field becomes an ordinary integral over the distance from the charge:

$$V = -\int_{\infty}^{r} E \cdot dr = -\int_{\infty}^{r} \left(\frac{kQ}{r^2}\right) dr$$

The constant k and the charge Q may both be taken outside the integral, since neither alters as the test charge moves:

$$V = -kQ\int_{\infty}^{r} \left(\frac{1}{r^2}\right) dr = -kQ\left[-\frac{1}{r}\right]_{\infty}^{r} = -kQ\left(-\frac{1}{r} - 0\right) = -kQ\left(-\frac{1}{r}\right)$$

The potential of a point charge is therefore:

$$\mathbf{V = \frac{kQ}{r}}$$

Read this formula closely, for it is among the most useful in the whole subject. The potential falls off as $\frac{1}{r}$, more gently than the field's $\frac{1}{r^2}$: halve your distance from the charge and the potential doubles, where the field would quadruple.

The sign of the potential follows the sign of the charge. Around a positive charge the potential is everywhere positive, sloping upward as the charge is approached, a hill; around a negative charge it is everywhere negative, a valley. And, decisively, $V = \frac{kQ}{r}$ carries no direction with it. It is a pure number at each point, a count of joules per coulomb.

Since k is the same constant that appears in Coulomb's law, the potential is just as often written as:

$$\mathbf{V = \frac{Q}{4\pi\varepsilon_0 r}}$$

Now to the promise. Suppose several point charges are present together, Q_1, Q_2, Q_3 and so on, lying at distances r_1, r_2, r_3 from some point of interest. *What is the potential there?* Because potential is energy per unit charge, and energies simply add, the potential of the whole group is the ordinary sum of the potentials each charge would produce alone:

$$\mathbf{V = \frac{kQ_1}{r_1} + \frac{kQ_2}{r_2} + \frac{kQ_3}{r_3} + \cdots}$$

This is the **superposition of potentials**, and it is worth pausing on how much lighter it is than the superposition of fields. To combine fields, every field's direction had to be tracked, the vectors resolved into components, the components added separately and the whole reassembled, with an angle to compute at the end. To combine potentials there are no directions and no components: each charge offers a single number, $\frac{kQ}{r}$, and the numbers are simply added. The one care needed is with signs, *for a negative charge contributes a negative potential, and that minus sign must be carried faithfully into the sum.* A positive and a negative charge at equal distances from a point contribute potentials that cancel each other exactly, leaving nothing.

An example will show the method at work, and the figure sets the scene: four equal positive charges, one pinned at each corner of a square, with the potential to be found at the dead centre.

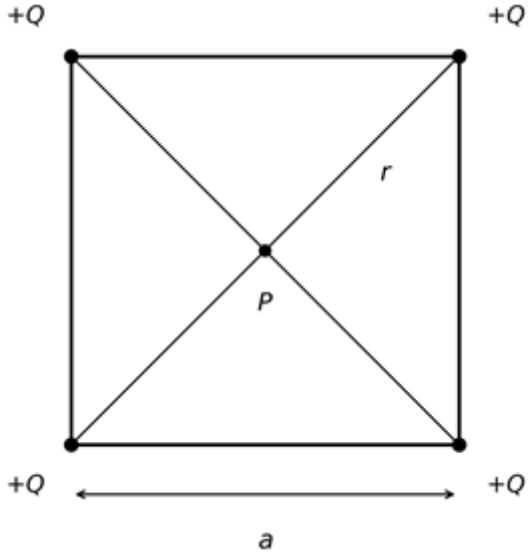


Figure: ***Four equal positive charges, one at each corner of a square of side a.*** *The centre P lies the same distance r from every corner, so each charge contributes an equal share to the potential there.*

The potential difference between two finite points

The result $V = \frac{kQ}{r}$ just settled is the **absolute** potential at a point distance r from a point charge Q, taken with the conventional zero of potential at infinity. In practice, however, an absolute potential is not the natural quantity. The thing that does the physical work, that drives charges through a wire and accelerates them in a field, is the **difference** of potential between two points. It is the difference that the equation $V_{AB} = V_A - V_B$, met earlier, defines, and it is the difference that a voltmeter actually measures. The absolute potential is a special case of the difference, recovered when one of the two points is sent to infinity.

For a point charge Q at the origin and two field points A and B at radial distances r_A and r_B from it, the potential difference follows directly from $V = kQ/r$ applied at each point:

$$\mathbf{V_{AB} = V_A - V_B = \frac{kQ}{r_A} - \frac{kQ}{r_B} = kQ\left(\frac{1}{r_A} - \frac{1}{r_B}\right)}$$

This is the work per unit positive charge that an external agent must do to carry a charge from B to A. It is finite for any two finite radial distances. Sending B to infinity, the reciprocal $\frac{1}{r_B}$ vanishes $\left(\frac{1}{\infty} \approx 0\right)$ and the equation collapses to the absolute potential $V_A = \frac{kQ}{r_A}$.

The absolute potential is therefore just a difference in which one of the two points has been pushed where its contribution can be safely ignored. The version with two finite points is the one that survives in distributions for which infinity is not a sensible reference, the infinite line of charge being the most important example.

Four charges, one tidy square, and a question that the scalar nature of potential turns from a chore into barely three lines of arithmetic.

BINDER Example 19

Four equal point charges, each of $8.0 \times 10^{-9}\text{C}$, are fixed one at each corner of a square of side 0.40m. Find the electric potential at the centre of the square. Take $k = 8.99 \times 10^{9}\text{Nm}^2/\text{C}^2$.

Solution

The centre of a square lies the same distance from all four of its corners. That distance is half the length of the diagonal, $r = \frac{a}{\sqrt{2}}$, which for a side of 0.40m comes to $r = \frac{0.40\text{m}}{\sqrt{2}} = 0.28\text{m}$.

Electric potential is a scalar, so the potential at the centre is simply the sum of the four potentials, one from each corner charge. The four charges are equal and all lie the same distance r from the centre, so the four contributions are identical, and the total is four times any one of them:

$$V = 4\frac{kQ}{r}$$

Substituting:

$$V = \frac{4 \times 8.99 \times 10^{9}\text{Nm}^2/\text{C}^2 \times 8.0 \times 10^{-9}\text{C}}{0.28\text{m}} = 1.0 \times 10^{3}\text{V}$$

The electric potential at the centre of the square is about $1.0 \times 10^{3}\text{V}$, that is one kilovolt.

Making Sense of the Answer: *It is worth setting this result beside what the electric field does at the very same point. The four charges are equal and sit symmetrically about the centre, so the four field vectors there are equal in length and point in four opposing directions; they cancel in perfect pairs, and the field at the centre is exactly zero. The potential at that same point is a thousand volts. Field nothing, potential everything, and no contradiction in it. The field is a vector, and four vectors set symmetrically can sum to nothing; the potential is a scalar, and four positive numbers can only pile up.*

Think Like a Physicist: *Let the scalar nature of potential shape how you plan a problem. Whenever several charges are in play, ask first what the question truly wants. If it asks for an energy, a work done, or a speed gained, it wants the potential, and the potential is a single number got by plain addition. If it insists on the field, on the direction a force will point, then the vectors cannot be dodged. But here is the move worth keeping: even when the field is what you finally want, it is often easier to find the potential everywhere first, as one scalar map laid over space, and then read the field off from the way that potential changes from place to place. The field is the slope of the potential, a relationship glimpsed in the opening section and soon to be made exact. Scalar first, vector afterwards, is for a great many problems the shorter road.*

Electric potential of a finite line of charge

Every potential met so far in this section has been built up from point charges. The same recipe extends bodily to a continuous distribution: cut the distribution into infinitesimal pieces, each effectively a point charge of charge dq, compute the potential each piece would contribute at the field point using the point-charge formula, and add the contributions by integral. As the cleanest illustration of the method, take a thin straight rod of length 2a, uniformly charged with linear density λ, and let P sit at perpendicular distance x from the midpoint of the rod, as the figure shows.

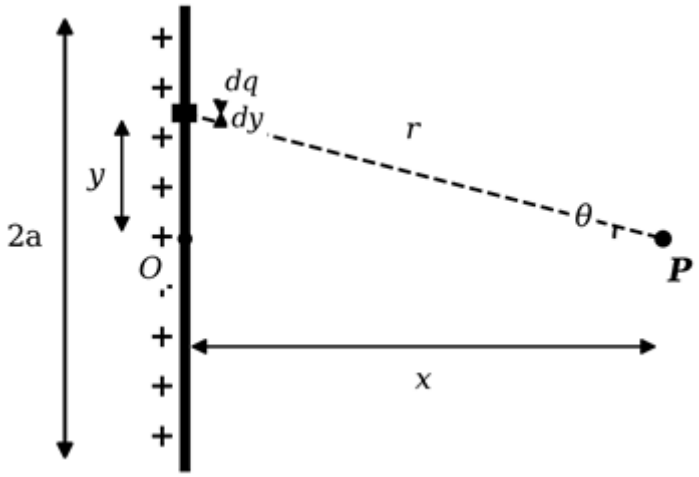


Figure: ***Geometry for the finite line of charge.*** *The same geometry used earlier. The element of length dy at height y carries charge dq = λ dy and sits at distance r from P. The potential it contributes at P is dV = k dq / r, a scalar that adds simply to the contributions of every other element along the rod.*

Each element of length dy along the rod carries the small charge $dq = \lambda\, dy$, and lies at distance $r = \sqrt{x^2 + y^2}$ from P. The potential this single element would contribute at P, treating it as a tiny point charge, is the point-charge formula applied to its charge:

$$dV = \frac{k\,dq}{r} = \frac{k\lambda\,dy}{\sqrt{x^2 + y^2}}$$

Adding the contributions of every element along the rod, with y running from $-a$ to a:

$$\mathbf{V} = k\lambda \int_{-a}^{a} \frac{dy}{\sqrt{x^2 + y^2}} = k\lambda \left[\ln\left(y + \sqrt{x^2 + y^2}\right)\right]_{-a}^{a} = \mathbf{k\lambda ln}\left(\frac{\sqrt{\mathbf{x^2 + a^2}} + \mathbf{a}}{\sqrt{\mathbf{x^2 + a^2}} - \mathbf{a}}\right)$$

Substituting $k = \frac{\lambda}{4\pi\varepsilon_0}$:

$$\mathbf{V} = \frac{\boldsymbol{\lambda}}{\mathbf{4\pi\varepsilon_0}} \mathbf{ln}\left(\frac{\sqrt{\mathbf{x^2 + a^2}} + \mathbf{a}}{\sqrt{\mathbf{x^2 + a^2}} - \mathbf{a}}\right)$$

This is the electric potential at perpendicular distance x from the midpoint of a uniformly charged rod of length 2a, taken with the usual reference of zero potential at infinity. Two short remarks close the case.

- The result depends only on the perpendicular distance from the midpoint, not on the orientation of the rod in space; the potential is a scalar, and the contributions of every element add by simple arithmetic, free of any worry about directions.
- The formula degenerates in an instructive way when the rod is allowed to grow infinitely long, a tending to infinity. The argument of the logarithm grows without bound, and the potential at any finite point P diverges.

From the finite line to its infinite-line difference

When the finite rod is allowed to grow infinitely long, a tending to infinity, the absolute potential of the previous derivation diverges, as already noted. The remedy is based on the fact that: the natural quantity is no longer the absolute potential but the difference of potentials between two points, and that difference is finite even when each of its terms, taken on its own, is not. The line of charge is the showcase example.

For an infinite straight line carrying linear charge density λ, the field at perpendicular distance r from the line is, as derived earlier,

$$E = \frac{\lambda}{2\pi\varepsilon_0 r}$$

Take two field points at perpendicular distances r_a and r_b from the line. The potential difference between them follows from the general relation $V = -\int \vec{E} \cdot d\vec{r}$, integrated radially from the second point to the first:

$$\mathbf{V_a - V_b} = -\int_{r_b}^{r_a} \frac{\lambda}{2\pi\varepsilon_0\, r}\, dr = -\frac{\lambda}{2\pi\varepsilon_0}[\ln r]_{r_b}^{r_a} = \frac{\boldsymbol{\lambda}}{\mathbf{2\pi\varepsilon_0}} \mathbf{ln}\left(\frac{\mathbf{r_b}}{\mathbf{r_a}}\right)$$

The result is perfectly finite for any two finite radii. The logarithm grows only slowly and never diverges away from the special value $r = 0$ of the line itself. The divergence of the absolute potential, which looked like a failure of the physics when it first appeared, was nothing of the sort; it was simply a sign that infinity was the wrong reference for this distribution, and any other finite reference will do. The line of charge is the prototype for every distribution of infinite extent.

The conducting cylinder, as a corollary

A natural choice of reference for the line of charge is the surface of the line itself, when the line is realised in practice as a long conducting cylinder of radius R, uniformly charged. Setting the potential at the cylinder's surface to zero, $V(R) = 0$, and applying the difference formula with $r_b = R$ and $r_a = r$ for any $r > R$:

$$V(r) = \frac{\lambda}{2\pi\varepsilon_0} \ln\left(\frac{R}{r}\right)$$

This is the potential at perpendicular distance r from the axis of a long charged conducting cylinder, taken with the cylinder's own surface as the zero of potential. The formula is negative for any $r > R$ (the logarithm of a fraction less than one is negative), which is the right sign: the potential falls as one moves away from a positively charged conductor, and falling below the zero set at the surface means becoming negative. Inside the conductor itself, the field is zero throughout, so the potential is constant; it must therefore take everywhere within the metal the value it had on the surface, which is the zero we chose.

Electric potential of a charged conducting sphere

One further continuous distribution is worth its turn before the dipole, and a familiar one. The conducting sphere whose field was settled in the previous section is also a continuous distribution of charge, and its potential, by the relation $V = -\int \vec{E} \cdot d\vec{r}$ that we have just been working with, follows from the field with no fresh effort. The result is worth seeing because it places side by side, in a single picture, the most useful contrast in all of electrostatics: a field that jumps abruptly at a surface and a potential that does not.

Take a solid conducting sphere of radius R, carrying a total positive charge Q, left in electrostatic equilibrium. The field, recalled from the previous section, is the field of a point charge Q sitting at the centre outside the sphere, and is zero inside the conductor:

$$E = \frac{kQ}{r^2} \quad \text{for} \quad r > R, \quad E = 0 \quad \text{for} \quad r \leq R$$

The potential at any point in space, taken with the conventional zero at infinity, follows from integrating this field along a path that runs in from infinity to the point in question. Two regions need separate attention.

For a point outside the sphere, at distance r from the centre with $r > R$, the integration is exactly the integration that was done for a point charge. The field along the inward path is kQ/r'^2, and

$$V(r) = -\int_{\infty}^{r} \frac{kQ}{r'^2}\, dr' = \frac{kQ}{r}$$

From outside, then, a conducting sphere looks like a point charge of total Q at its centre, just as far as the potential is concerned as it does as far as the field is concerned. At the surface itself, $r = R$, this gives:

$$V(R) = \frac{kQ}{R}$$

For a point inside the conductor, $r < R$, the field is zero throughout. The change in potential between any two points inside the conductor is therefore zero as well, since

$$dV = -\vec{E} \cdot d\vec{r} = 0 \quad \text{inside.}$$

The potential cannot change anywhere within the body of the conductor, and must take everywhere the value it had at the surface. The whole interior of the conductor sits at the constant potential.

$$V = \frac{kQ}{R} \quad \text{for} \quad r \leq R$$

Drawing the two regions together: the potential of a charged conducting sphere is constant at the value $\frac{kQ}{R}$ throughout its whole interior, falls smoothly as kQ/r outside, and joins continuously from one regime to the other at the surface $r = R$. The figure draws the picture, and sets the potential alongside the field for direct comparison.

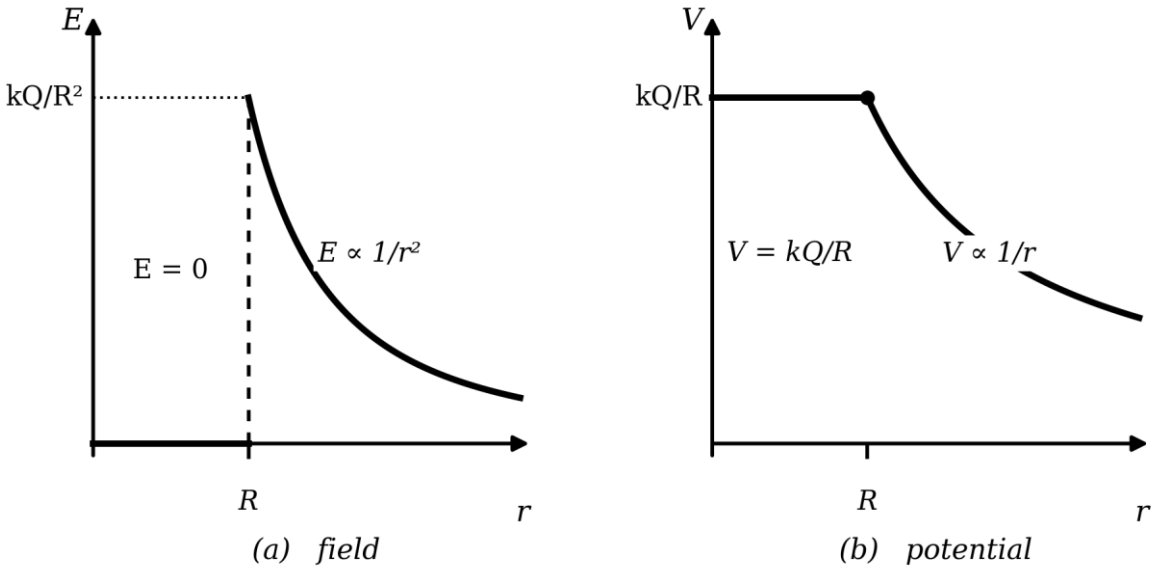


Figure: ***Electric field and electric potential of a charged conducting sphere.*** *Panel (a): the field is zero inside the conductor, jumps abruptly to its peak value kQ/R² at the surface, and falls as the inverse square outside. Panel (b): the potential is constant at kQ/R throughout the interior, joins continuously to the curve kQ/r at the surface, and falls as the inverse first power outside. The field is discontinuous at the surface; the potential is not.*

The two graphs are worth looking at together, because they tell the same story in two different scripts.

- Outside the sphere both quantities fall away with distance, the field as the inverse square and the potential as the inverse first power, the slower fall reflecting the fact that the potential is the running sum of the field's effect inward from infinity.
- At the surface itself the two part company in a way worth noticing. The field is discontinuous there: it leaps from zero just inside to its peak value kQ/R^2 just outside, a jump that records the thin layer of charge mathematically present on the idealised conductor's surface. The potential, by contrast, is perfectly continuous at the surface: the constant $\frac{kQ}{R}$ carried out from the interior matches exactly the value the outside formula kQ/r delivers at $r = R$, and the curve sails through the surface without a kink. The contrast is precisely the one that the next section will name: the field is the slope of the potential, and a potential with a finite step in its slope answers with a finite jump in the field.

This is also the formal underpinning of the puzzle that runs through the next section. Inside a hollow charged conductor, met earlier in the catalogue of Gauss applications, the field is zero but the potential is emphatically not. The conducting cavity sits

at the constant surface value kQ/R_2, a region of zero field at non-zero potential that finds its cleanest statement in the present formula and its cleanest picture in the present figure. Field and potential, the chapter is beginning to say, are independent enough that either can vanish without the other doing the same, and a careful student keeps both in view at every step.

So far the charges added have all carried the same sign, and four positive numbers can do nothing but pile up. The next arrangement is the most important of all the simple ones, and the first to set a positive charge against a negative one: a single pair, plus and minus, held a small fixed distance apart. It is called the electric dipole, and superposing its two potentials, one positive and one negative, yields a result whose shape is surprising, widely useful, and the perfect illustration of how a field and a potential can part ways.

The electric dipole

The previous section closed on a promise of the most instructive of all the simple charge arrangements. This section makes good on it.

A positive charge and a negative charge, equal in magnitude and held a small fixed distance apart: this modest pair is called the **electric dipole**. Modest it may look, but it is one of the most important arrangements in all of physics. Every water molecule is a tiny dipole, its oxygen end slightly negative and its hydrogen end slightly positive, and that single fact is why water dissolves so much and behaves so strangely. Radio aerials are driven dipoles. An atom placed in an electric field stretches into a dipole. The arrangement repays close study.

Two numbers fix a dipole: the **size of the charge, q,** and the **separation between the charges, d.** The two are bundled into a single quantity, the **dipole moment**, written **p** and defined as *the product of either charge and the separation vector between the charges in an electric dipole*.

Mathematically:

$$\mathbf{p} = \mathbf{q}\vec{\mathbf{d}}$$

The dipole moment is a vector: it points along the line joining the charges, *from the negative charge towards the positive*, and its size qd measures how strong the dipole is. Everything a dipole does at a distance, it does through its dipole moment p.

Two questions now press: *what potential does a dipole set up in the space around it, and what field?* We shall answer both, at two telling places: along the dipole's **axis**, the line through both charges, and along its **perpendicular bisector**, the line that cuts the dipole in half at right angles. The figure marks them.

A point A sits on the axis and a point B on the bisector, each a distance r from the centre of the dipole. Throughout, we take r to be large compared with the small separation d, which is the case of real interest.

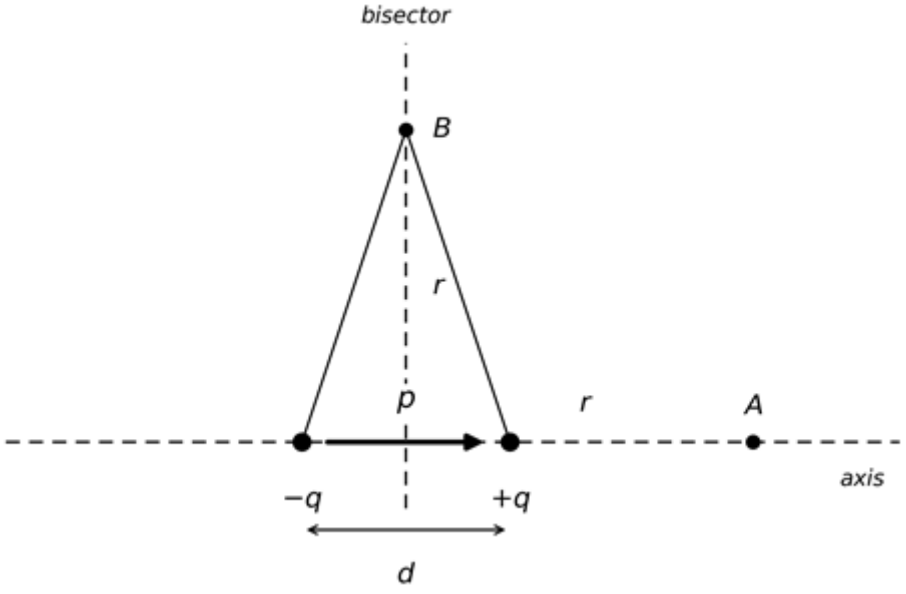


Figure: *An electric dipole: charges $+q$ and $-q$ a distance d apart, with dipole moment p pointing from the negative charge to the positive. The point A lies on the axis and B on the perpendicular bisector, each a distance r from the centre.*

Begin with the potential at A, on the axis. The point stands a distance $r - \frac{d}{2}$ from the nearer charge, the positive one, and $r + \frac{d}{2}$ from the farther charge, the negative one. Potential is a scalar, so the potential at A is the plain sum of the two contributions, each carrying the sign of its charge:

$$V = \frac{kq}{\left(r - \frac{d}{2}\right)} - \frac{kq}{\left(r + \frac{d}{2}\right)} = \frac{kqd}{\left(r - \frac{d}{2}\right)\left(r + \frac{d}{2}\right)} = \frac{kqd}{r^2 - \frac{d^2}{4}}$$

Now the case of interest. When A is far from the dipole compared with its tiny separation, r much greater than d, the term $\frac{d^2}{4}$ is negligible beside r^2 and may be dropped. Writing the dipole moment $p = qd$ in place of qd, the potential on the axis of a dipole is:

$$\mathbf{V = \frac{kqd}{r^2} = \frac{kp}{r^2}}$$

Notice the inverse-square falloff. A single point charge's potential dies away as $\frac{1}{r}$; a dipole's, on its axis, dies faster, as $\frac{1}{r^2}$. The reason is the dipole's neutrality: from far off its two charges very nearly cancel, and only their slight separation keeps any potential alive at all.

Now the bisector, and the point B. Here something different happens. B lies on the line that cuts the dipole symmetrically in half, so it is exactly the same distance from the positive charge as from the negative one. Call that common distance s. The positive charge raises the potential at B by $\frac{kq}{s}$; the negative charge lowers it by precisely the same $\frac{kq}{s}$; and the two cancel without remainder:

$$\mathbf{V = \frac{kq}{s} - \frac{kq}{s} = 0}$$

This zero is exact, and no approximation was made. It holds not at one special point but at every point along the whole perpendicular bisector, near the dipole or far from it. The bisector is a line of zero potential.

So much for the potential. We turn now to the field, and we shall let the bisector spring its surprise in a moment.

First the axis. The field at A is found by the same superposition, but now of the two fields rather than the two potentials, and a field is a vector. Both fields here lie along the axis: the positive charge pushes outward, away from the dipole, with magnitude $\frac{kq}{(r-\frac{d}{2})^2}$; the negative charge pulls inward, back towards itself, with magnitude $\frac{kq}{(r+\frac{d}{2})^2}$. Taking the outward direction as positive, the net field is their difference:

$$E = \frac{kq}{\left(r-\frac{d}{2}\right)^2} - \frac{kq}{\left(r+\frac{d}{2}\right)^2} = \frac{kq\left((r+\frac{d}{2})^2 - (r-\frac{d}{2})^2\right)}{\left(r-\frac{d}{2}\right)^2\left(r+\frac{d}{2}\right)^2} = \frac{2kqrd}{\left(r^2-\frac{d^2}{4}\right)^2}$$

Once more taking r much larger than d, the denominator becomes simply r^4, and with $p = qd$ the field on the axis of a dipole is:

$$\mathbf{E = \frac{2kqd}{r^3} = \frac{2kp}{r^3}}$$

The falloff is faster still, as $\frac{1}{r^3}$, and the field points the same way as the dipole moment p.

And now the bisector. The potential there was zero; *is the field zero too?* It is not. At B each charge produces a field of magnitude $\frac{kq}{s^2}$, the positive charge's pointing away from it and the negative charge's pointing towards it. Resolve each of these two fields into a part along the axis and a part along the bisector. The parts along the bisector are equal in size and opposite in direction, one up and one down, and cancel exactly. The parts along the axis are equal in size and point the same way, and so they add. Each axial part is the field magnitude $\frac{kq}{s^2}$ times the cosine of its angle to the axis, which the geometry fixes at $\frac{d}{2s}$. Two equal axial parts therefore give, with $s = \sqrt{r^2 + \frac{d^2}{4}}$:

$$E = 2 \times \frac{kq}{s^2} \times \left(\frac{d}{2s}\right) = \frac{kqd}{s^3}$$

With r much greater than d the distance s becomes r, and with $p = qd$ the field on the bisector of a dipole is:

$$\mathbf{E = \frac{kqd}{r^3} = \frac{kp}{r^3}}$$

It points along the axis but antiparallel to the dipole moment; it is therefore opposite to the field on the axis, and exactly half that size at the same distance.

Set the two bisector results side by side, for together they carry the lesson this section has been driving towards.

Along the perpendicular bisector of a dipole the potential is zero, exactly and everywhere. Along that very same line the electric field is $\frac{kp}{r^3}$, and most certainly not zero. A student who argues that the potential here is zero and so the field here must be zero has made a real and tempting mistake. Potential is energy per unit charge; field is force per unit charge; they are different quantities, answering different questions, and a place may be empty of the one while full of the other. The bisector of a dipole is the plainest proof of it in all of electrostatics: zero potential, and a living field, on the very same line.

A short calculation will fix the zero-potential result in the mind, by checking it at the one point easiest to reach: the dipole's own midpoint.

BINDER Example 20

Two point charges, $+3.0 \times 10^{-9}\text{C}$ and $-3.0 \times 10^{-9}\text{C}$, are fixed 20cm apart. Find the electric potential at the midpoint of the line joining them. Take $\text{k} = 8.99 \times 10^{9}\text{N} \cdot \text{m}^2/\text{C}^2$.

Solution

The midpoint lies halfway between the charges, a distance of 10cm, or 0.10m, from each. Potential is a scalar, so the potential at the midpoint is the sum of the two contributions, one from each charge, each carrying the sign of its charge:

$$\text{V} = \frac{\text{k}(+\text{q})}{\text{r}} + \frac{\text{k}(-\text{q})}{\text{r}}$$

Both charges lie the same distance from the midpoint and are equal in magnitude, so the two contributions are equal in size; but one charge is positive and the other negative, so the two are opposite in sign. Substituting $3.0 \times 10^{-9}\text{C}$ and $\text{r} = 0.10\text{m}$:

$$\text{V} = \frac{\text{k}(3.0 \times 10^{-9}\text{C})}{0.10\text{m}} - \frac{\text{k}(3.0 \times 10^{-9}\text{C})}{0.10\text{m}} = 0$$

The potential at the midpoint is zero.

Making Sense of the Answer: *The zero is worth dwelling on. It did not come about because the charges are far away or weak; each charge, on its own, sets up a healthy 270 volts at the midpoint. It came about because the two contributions are equal and opposite, and a scalar sum lets them cancel completely. And the midpoint is nothing special: it is simply the most convenient point on the perpendicular bisector at which to test. Every other point on that bisector would have given zero by the identical argument, the two charges always equidistant and their potentials always equal and opposite. This single calculation stands for a whole line of zeros. Yet, as the section showed, the electric field marches straight through that line of zero potential without itself vanishing at all.*

Think Like a Physicist: *Carry one habit away from the dipole: keep the sign of every charge with it, always, into every potential sum. The whole result here, the clean cancellation to zero, lives or dies on the minus sign of the negative charge. Drop it, treat both contributions as positive, and the midpoint potential comes out as 540 volts instead of nothing: a wrong answer produced by a careless hand rather than a wrong method. Potential is a scalar, which makes it easy; but a scalar still carries a sign, and the sign is not optional. When charges of both kinds are present, write each contribution as the charge, complete with its sign, divided by its distance and multiplied by k, and let the arithmetic take care of the rest.*

The dipole has shown how potentials combine, signs and all, and how sharply a dipole's reach falls away with distance. One large question about potential now remains, and it is the one that ties this whole section back to energy. We have spoken throughout of the work done in moving a charge, and of energy carried per unit charge; it is time to speak of the energy itself: the electric potential energy stored when charges are assembled against their will, and released when they are let fly apart. That is the business of the section's final part.

Electric potential energy

The section's final task is the one its very first definition pointed towards. Potential was introduced as energy carried per unit charge; it remains now to speak of that energy itself.

Every charge that sits in an electric field holds a store of energy by virtue of where it sits, and that store is its **electric potential energy**. The potential introduced at the start of this section delivers it at once. Potential was defined as energy per unit charge, which is to say that the potential V at a point is the potential energy belonging to each coulomb of charge placed there. A charge q is q coulombs, so its potential energy is q times the energy carried by each:

$$\mathbf{U = qV}$$

The unit is the joule, as a charge in coulombs multiplied by a potential in volts is bound to give. The sign looks after itself. A positive charge at a point of positive potential has positive energy, the signature of a charge pushed into place against a repulsion, which would spring back the moment it were freed. A positive charge at a point of negative potential has negative energy, the signature of a charge that is bound, one that could be dragged away to infinity only if work were done on it from outside.

The most important case is also among the simplest. It is not one charge sitting in another's field, but two charges together, and the energy that belongs to the pair of them. Take two point charges, Q_1 and Q_2, a distance r apart, and ask how much energy is stored in that arrangement. The honest way to count it is to ask how much work was needed to build the arrangement at all, starting with the two charges infinitely far apart.

Bring in Q_1 first. Empty space resists it not at all, so placing the first charge costs no work. With Q_1 in position it has spread a potential through the space around it, and at the distance r where Q_2 is to go that potential has the value $\frac{\text{kQ}_1}{\text{r}}$. Carrying Q_2 in to that spot costs work equal to Q_2 multiplied by the potential already waiting there, and that work is the potential energy of the pair:

$$\mathbf{U = Q_2 \times \frac{kQ_1}{r} = \frac{kQ_1Q_2}{r}}$$

The result makes no distinction between the two charges. Bringing Q_1 in to the potential of Q_2 instead would have produced the identical $\frac{kQ_1Q_2}{r}$; the energy belongs to the pair, not to either charge by itself. Its sign tells the pair's whole story.

Two like charges, both positive or both negative, give a positive product Q_1Q_2 and so a positive energy: work was spent forcing them together against their mutual repulsion, and that work is stored, ready to be handed back the instant they are released. Two unlike charges give a negative product and a negative energy: the pair is bound, and energy would have to be supplied from outside before they could be parted to infinity. Positive energy is a coiled spring; negative energy is a knot.

Three charges, or four, or more, bring no new principle. The energy of any collection of point charges is the sum of the energies of all the separate pairs that can be drawn within it. For three charges (Q_1, Q_2, Q_3) there are three pairs, and the total energy is the three pair-energies added together:

$$U = \frac{kQ_1Q_2}{r_{12}} + \frac{kQ_1Q_3}{r_{13}} + \frac{kQ_2Q_3}{r_{23}}$$

Each term is one pair's $\frac{kQ_1Q_2}{r}$, carrying its own two charges and its own separation. The figure shows the three pairs of a triangle of charges. Four charges would furnish six pairs, five charges ten, and the recipe never alters: find every distinct pair, price it with $\frac{kQ_iQ_j}{r}$, and add.

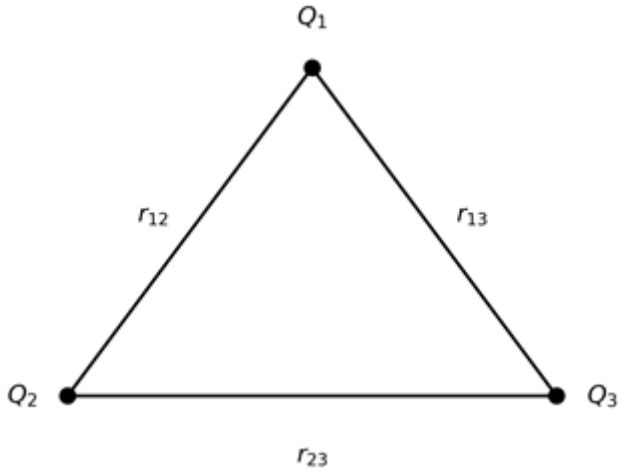


Figure: ***Three point charges form three pairs.*** *The energy of the system is the sum of the three pair-energies, one for each side of the triangle.*

One caution attaches to this counting, and it is the source of a factor of one-half that often appears alongside the formula. The pair-by-pair sum just described counts each pair exactly once, and is correct as it stands. A tempting alternative is to visit each charge in turn, find the potential at its location due to all the others, multiply by the charge itself to obtain its qV, and add those products. That alternative counts every pair twice, once at each of the two charges that form it. To undo the double-counting the whole sum must be halved, and that halving is the one-half. The pair-by-pair count needs no such repair, and is the safer of the two.

A last and very practical question closes the section: *how much energy does it take to charge a conductor, to gather a total charge Q onto an isolated metal sphere of radius R?*

The charge cannot arrive all at once. It must be carried on in small instalments, and each instalment proves harder to add than the one before. A sphere, seen from its surface or from beyond, behaves exactly as a point charge at its centre would, so while it carries a charge q its surface rests at the potential $\frac{kq}{R}$. Adding the next small instalment dq means pushing it onto a sphere already held at that potential, and the work for that one small step is the potential multiplied by the small charge, $dW = V\,dq$. The total work is the sum of every such step as the charge builds from nothing to its final value Q, which is an integral:

$$W = \int_0^Q V\,dq = \int_0^Q \left(\frac{kq}{R}\right) dq$$

The constant k and the radius R both come outside the integral, and the integral of q is $\frac{q^2}{2}$:

$$W = \left(\frac{k}{R}\right)\int_0^Q q\,dq = \left(\frac{k}{R}\right)\left[\frac{q^2}{2}\right]_0^Q$$

Evaluating the bracket between its limits, $\frac{q^2}{2}$ is $\frac{Q^2}{2}$ at the upper limit and zero at the lower, and the energy stored in charging the conductor is:

$$\mathbf{W = \frac{kQ^2}{2R}}$$

Read what the integral has done. The stored energy grows not as the charge Q but as its square. Double the charge carried by the sphere and four times the energy is stored, not twice. The reason lies in that climbing potential: the first faint instalments of charge are added almost for nothing, onto a sphere of low potential, while the final instalments must be forced onto a sphere whose potential has risen high. The average cost across the whole process comes to exactly half the final potential, and that one-half is why the energy is $\frac{kQ^2}{2R}$ and not $\frac{kQ^2}{R}$. The same one-half can be seen directly: the result is one-half of the final charge Q multiplied by the final surface potential $\frac{kQ}{R}$.

A short and concrete calculation will set the energy of a pair to work.

BINDER Example 21

Two point charges, $+5.0 \times 10^{-9}\text{C}$ and $+2.0 \times 10^{-9}\text{C}$, are held 30cm apart. Find the electric potential energy of the pair. Take $k = 8.99 \times 10^9 \text{Nm}^2/\text{C}^2$.

Solution

The potential energy of a pair of point charges is $\frac{kQ_1Q_2}{r}$, each charge entered complete with its sign and r the separation between them:

$$U = \frac{kQ_1Q_2}{r}$$

Both charges are positive, and the separation of 30cm is 0.30m. Substituting, with every quantity in SI units:

$$U = \frac{8.99 \times 10^9 \text{Nm}^2/\text{C}^2 \times 5.0 \times 10^{-9}\text{C} \times 2.0 \times 10^{-9}\text{C}}{0.30\text{m}} = 3.0 \times 10^{-7}\text{J}$$

The potential energy of the pair is about $3.0 \times 10^{-7}\text{J}$.

Making Sense of the Answer: *The energy is positive, and that positive sign is the physics of the answer. Both charges are positive, so they repel; to have brought them as close together as 30 centimetres, work had to be spent against that repulsion, and the result is precisely that spent work, now held in the arrangement. Release the charges and the store is repaid in full: they fly apart, and the energy reappears as kinetic energy. Had one charge been negative, the product of the two charges would have been negative, the energy negative with it, and the meaning turned around: a bound pair, holding together, needing energy supplied before it could ever be separated.*

Think Like a Physicist: *Two habits from this section are worth carrying onward. The first is that potential energy, like potential itself, is a scalar, and it is counted pair by pair. When a problem sets several charges before you, do not hunt for a single clever formula; find every pair, price each with the pair formula, and add the prices, signs and all. The second is to watch for the square. The energy stored in an arrangement of charge very often grows as the square of that charge, because the act of building the charge up raises the very potential it must be built against. Whenever doubling a charge is found to quadruple an energy, the integral met in this section is quietly at work behind the result.*

With this the section on electric potential and potential energy is complete. Potential difference, the potential of a single point charge, the potential and field of a dipole, and the energy stored in assembled charge have all been set out and worked through. What follows is a gathering of further worked examples, the real-world cases and the harder ones, drawing on everything the section has built.

Further worked examples on electric potential and potential difference

The four parts behind us built the tools; the examples that follow are where the tools are made to work. The physics does not change in them. The questions simply grow more demanding, and one or two will look, at first sight, as though nature has broken its own rules. It has not.

REAL Example 22

A bird settles on a single bare conductor of a high-voltage transmission line and perches there, both feet gripping the one wire, entirely unharmed, even though that wire may be carrying electricity at many thousands of volts. A person who took hold of the same wire while standing on the ground would be killed at once. Explain, in terms of electric potential, why the bird is safe.

Solution

Whether a current is driven through a body depends not on the potential of that body but on the potential difference between its parts. The potential difference between two points is the work needed, per unit charge, to carry charge from one to the other; where there is no potential difference, no work is done on the charge, and nothing is there to drive a current.

The bird grips the wire with two feet set only a few centimetres apart. The wire is an excellent conductor, and along so short a length of good conductor there is almost no potential difference at all: the bird's two feet rest at the same potential, whatever

great value that potential may hold relative to the distant ground. With no potential difference between its feet, no work is done on the charges in the bird's body, no current flows through it, and the bird feels nothing.

The bird is safe not because the wire is at a low potential, for it is not, but because both of its feet touch the wire at one and the same potential. The person on the ground stands in a wholly different case. One hand grips the wire at a potential of thousands of volts while the feet rest on the earth at a potential near zero. The body then bridges an enormous potential difference, a large current is driven through it from hand to foot, and the outcome is fatal. In electricity the danger lives in the potential difference, never in the potential alone.

REAL Example 23

The safety advice for anyone caught in the open during a fierce thunderstorm, or who comes upon a power line lying broken on the ground, is to crouch low with the two feet held tightly together, and on no account to walk normally or to run. Using the idea of electric potential, explain why feet held together are safer than feet set apart.

Solution

When a large current pours into the earth at a single place, where lightning strikes the ground or where a live wire falls upon it, that point of entry is raised to a very high potential. The earth is not a perfect conductor, so the potential does not stay high everywhere. It falls away steeply with distance from the point of entry, leaving the ground around that point carrying a slope of potential: high close to the entry, lower and lower further off.

A person standing on that sloping ground with the feet apart has each foot resting at a different potential, since each foot stands a different distance from the point of entry. The difference between those two potentials appears across the person's legs, and if it is large it drives a dangerous current up one leg and down the other.

Crouching with the feet pressed together takes the danger away. Two feet in nearly the same place rest at nearly the same potential; the potential difference between them falls almost to zero, and with no potential difference to drive it, little current passes through the body. The warning never to run carries the same physics: a long running stride sets the two feet far apart on the slope of potential, across a large potential difference.

This is the bird on the wire seen from the other side. The bird is safe because both its feet share a single potential, while the person is endangered because their feet, set apart on sloping ground, do not.

BINDER Example 24

Inside a uranium-238 nucleus two protons may sit a distance of about $6.0 \times 10^{-15}\text{m}$ apart. Treating each proton as a point charge of $+1.6 \times 10^{-19}\text{C}$, calculate the electric potential energy of this pair of protons. Take k = $8.99 \times 10^{9}\text{Nm}^2/\text{C}^2$.

Solution

The electric potential energy of a pair of point charges is the product of the constant k with the two charges, divided by their separation:

$$\text{U} = \frac{\text{kQ}_1\text{Q}_2}{\text{r}}$$

Both charges are protons, so each is $1.6 \times 10^{-19}\text{C}$, and the separation is $6.0 \times 10^{-15}\text{m}$. Substituting:

$$\text{U} = \frac{8.99 \times 10^{9}\text{N} \cdot \text{m}^2/\text{C}^2 \times (1.6 \times 10^{-19}\text{C})^2}{6.0 \times 10^{-15}\text{m}} = 3.8 \times 10^{-14}\text{J}$$

The electric potential energy of the pair of protons is about $3.8 \times 10^{-14}\text{J}$.

HOT Example 25

Four equal positive point charges, each of magnitude q, are carried in from infinity and fixed one at each corner of a square of side a. Find the total work that must be done to assemble them.

Solution

The total work done is equal to the electric potential energy of the finished arrangement.

The potential energy of a system of point charges is the sum of $\frac{\text{kQ}_1\text{Q}_2}{\text{r}}$ taken over every distinct pair of charges. Four charges form six distinct pairs. As the figure shows, four of these pairs lie along the sides of the square, separated by the side length a, and the remaining two lie along the diagonals, separated by the diagonal length $\text{a}\sqrt{2}$.

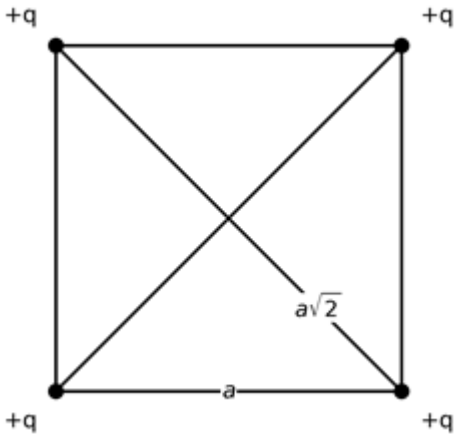


Figure: *Four charges at the corners of a square form six pairs: four pairs lie along the sides of the square and two along its diagonals.*

Each of the four side-pairs contributes $\frac{kq^2}{a}$, the two charges of a pair being both positive and equal.

Each of the two diagonal-pairs contributes $\frac{kq^2}{a\sqrt{2}}$.

The total energy is the sum of all six pairs:

$$U = 4 \times \left(\frac{kq^2}{a}\right) + 2 \times \left(\frac{kq^2}{a\sqrt{2}}\right)$$

The two diagonal terms simplify, since 2 divided by $\sqrt{2}$ is itself $\sqrt{2}$:

$$U = \left(4 + \sqrt{2}\right)\left(\frac{kq^2}{a}\right) \approx 5.41\left(\frac{kq^2}{a}\right)$$

HOT Example 26

A speck of dust of mass $5.0\mu g$ carries a charge of $2.0\mu C$. It is released from rest and moves freely, acted on by no force but the electric one, between two points whose potential difference is $2700V$. Find the speed the dust speck reaches.

Solution

As the charged speck moves through the potential difference, the electric force does work on it. The work done on a charge q carried through a potential difference ΔV is $q\Delta V$, and since no other force acts, all of that work becomes kinetic energy. The speck starts from rest, so its whole kinetic energy at the end is $\frac{1}{2}mv^2$:

$$q\Delta V = \frac{1}{2}mv^2$$

The speed v is wanted, so it is made the subject:

$$v = \sqrt{\frac{2q\Delta V}{m}}$$

Now the numbers, each in SI units: the charge $2.0\mu C$ is $2.0 \times 10^{-6}C$, the mass $5.0\mu g$ is $5.0 \times 10^{-9}kg$, and the potential difference is $2700V$:

$$v = \sqrt{\frac{2 \times 2.0 \times 10^{-6}C \times 2700V}{5.0 \times 10^{-9}kg}} = 1.5 \times 10^3 m/s$$

The dust speck reaches a speed of about $1.5 \times 10^3 m/s$.

HOT Example 27

A charge of $+3.0nC$ and a charge of $-6.0nC$ are fixed $24cm$ apart. Find every point on the straight line through the two charges at which the total electric potential is zero.

Solution

The potential at any point is the sum of the contributions of the two charges, each given by $\frac{kQ}{r}$ and each carrying the sign of its own charge. The positive charge raises the potential; the negative charge lowers it. The total is zero only where the positive contribution of the $+3.0nC$ charge is exactly matched by the negative contribution of the $-6.0nC$ charge.

Writing r_1 for the distance from the $+3.0nC$ charge and r_2 for the distance from the $-6.0nC$ charge, this balance is:

$$\frac{k(3.0nC)}{r_1} = \frac{k(6.0nC)}{r_2}$$

The constant k and the unit nanocoulomb stand on both sides and cancel, and the balance reduces to a plain statement about the two distances:

$$r_2 = 2r_1$$

Wherever the potential vanishes, the point must lie twice as far from the $-6.0nC$ charge as from the $+3.0nC$ charge. Two places on the line meet this condition. The first lies between the charges. There the two distances together make up the whole separation, so $r_1 + r_2 = 24cm$; with $r_2 = 2r_1$ this becomes $3r_1 = 24cm$, giving $r_1 = 8.0cm$. **One point of zero potential lies between the charges, $8.0cm$ from the $+3.0nC$ charge.**

The second lies outside the pair, beyond the $+3.0nC$ charge on the side away from the $-6.0nC$ charge. There the greater distance exceeds the smaller by the whole separation, so $r_2 - r_1 = 24cm$; with $r_2 = 2r_1$ this becomes $r_1 = 24cm$. **The second point of zero potential lies $24cm$ from the $+3.0nC$ charge, on its far side.**

Beyond the $-6.0nC$ charge there is no third point: that charge is the larger and would always be the nearer, so its negative potential can never be cancelled there.

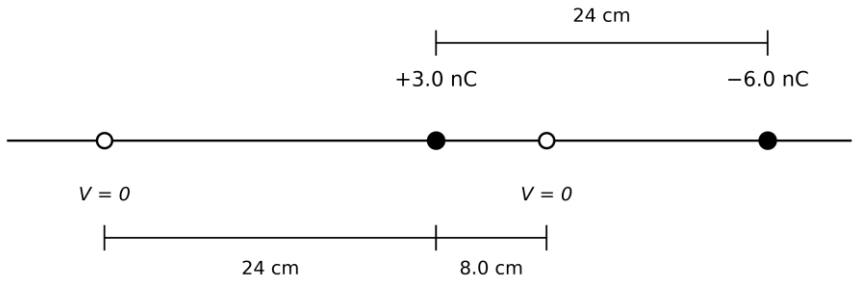


Figure: ***The two points of zero potential on the line of the charges:*** *one between the charges, the other beyond the positive charge.*

Make Sense of the Answer: *Two points of zero potential, and it is worth ending on what they are not. At neither of them is the electric field zero. Between the charges the field of the positive charge and the field of the negative charge point the same way, and two vectors in the same direction cannot cancel. At the outer point the two fields do oppose each other, but the point stands at unequal distances from two unequal charges, so their strengths differ and again they fail to cancel. Potential is a scalar: signed numbers cancel by simple addition, and zeros of potential are common. The field is a vector: it vanishes only where strength and direction conspire, at some quite different place. The lesson the dipole first taught holds to the end of the section. A point where the potential is zero is not a point where nothing is going on.*

From a bird on a wire to two protons inside a nucleus, these examples leaned on one small set of ideas: the potential of a point charge, the potential difference that alone drives a current, the energy bound up in a pair of charges and in a system of them, and the conservation of energy as a charge crosses a potential difference. Each question gave way to one of them. With this the section on electric potential and potential difference is complete, its formulas derived, its meaning drawn out, and its reach tested against the world.

The next section takes that potential one step further: it joins the points of equal potential into surfaces of their own, and shows that the electric field can be read straight from the way the potential changes from place to place.

EQUIPOTENTIAL SURFACES AND THE FIELD-POTENTIAL RELATIONSHIP

Anyone who has ever used a map of a mountainous place has already met equipotential surfaces, although under another name. On such a map the hills are drawn with thin closed lines that loop around one another, and every one of those lines is a contour: it joins together all the points on the ground that stand at the same height above the sea. Walk along a single contour line and a curious thing happens, you go neither up nor down. You may walk north for a while, then swing east, then curl back towards the south, but as long as your feet keep to the one line your height above the sea does not change by a single metre. To gain height you must do the very thing the contour forbids, you must step across it and onto the next line.

The map carries a second message, and it is hidden in the spacing of the lines. Where the contours are drawn far apart the ground rises gently, a long and easy walk between one marked height and the next. Where the contours crowd together into a tight band the ground is steep, the same gain in height now squeezed into a few hard strides. A walker who has learned to read a map never stops to measure a slope; the walker simply looks at how closely the lines run and knows at once where the climbing will be punishing and where it will be kind. The steepness of the hill is written in the spacing of its contours.

The electric potential, built up point by point in the previous section, possesses contour lines of exactly this kind. Gather together all the points in a field that happen to share one single value of potential, and they do not form a scattered dust of

separate points; they join up into a continuous surface. That surface is called an **equipotential surface,** and it is the electrical twin of the contour line on the mountain map. Carry a charge along it and, as the pages ahead will show, no work is done, just as the walker on the contour gains no height. Watch how the equipotentials of a field are spaced and the strength of the field can be read straight off them, just as the walker reads the steepness of the hill from the spacing of the lines. This section follows that one picture wherever it leads, and it leads a surprisingly long way.

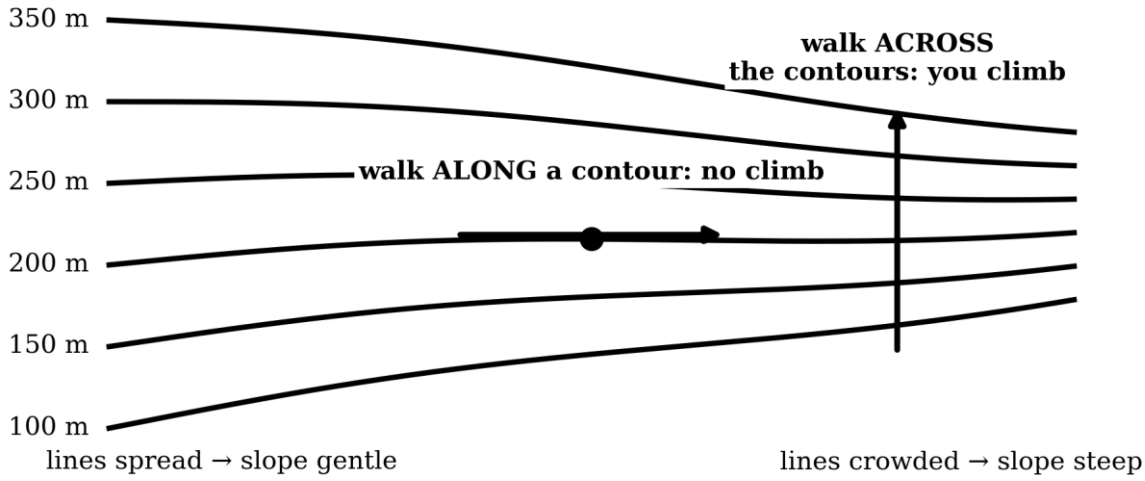


Figure: ***The contour-map analogy.*** *On a map of a hill each line joins points that stand at the same height above the sea. Walking along a single line involves no climbing at all, while the spacing of the lines reveals how steep the ground is: lines crowded close together mark steep ground, lines spread far apart mark gentle ground.*

Here, then, is the idea stated plainly. An **equipotential surface** *is a surface on which the electric potential has the same value at every point.* Move a test charge anywhere over such a surface and the potential it sits in never alters. Different surfaces carry different values, one for twelve volts, another for twenty, another for a hundred, and the whole family of them fills the space around a charge in the same way the contour lines fill a map, each surface a sheet of constant potential nested inside or outside its neighbours.

The shape of these surfaces is not left to guesswork; it follows at once from the formula for the potential. Around a single point charge the potential is $V = \frac{kQ}{r}$, a quantity that depends on nothing whatever except the distance from the charge. Every point lying the same distance from the charge therefore carries the same potential. And the set of all points at one fixed distance from a central point is, by its very definition, a sphere. The equipotential surfaces of a point charge are thus a family of spheres, all centred on the charge, one nested inside the next like the skins of an onion. On the flat of a page they can only be drawn as circles, but the reader should always picture the full spheres standing behind the circles.

Between two large parallel charged plates the picture is different but just as simple. There the field is uniform, and the potential climbs at a perfectly steady rate as one crosses the gap from the lower plate to the higher one. Every point the same distance from a plate sits at the same potential, and the set of all points at one fixed distance from a flat plate is itself a flat plane. The **equipotentials of the uniform field** are therefore a *stack of flat planes, all parallel to the plates and evenly spaced through the gap*. The two cases, the nest of spheres and the stack of planes, are drawn together in the figure.

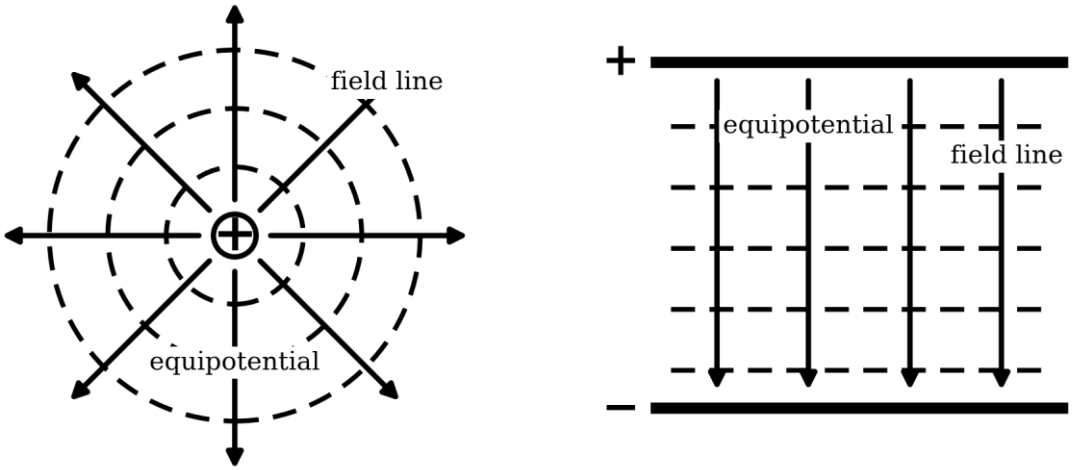


Figure: ***Equipotential surfaces.*** *Around a point charge the equipotentials are concentric spheres, drawn here as circles, and the field lines run straight out along the radii. Between parallel plates the equipotentials are flat parallel planes evenly spaced through the gap. In both cases the field crosses every equipotential at right angles.*

Equipotential surfaces are useful because a small handful of plain facts about them can be stated once and then leaned upon for the rest of electrostatics. Five such facts deserve to be set out carefully, each together with the reasoning that makes it true, for a fact whose reason is understood is a fact that can never afterwards be forgotten.

The first property is that ***the electric field is everywhere at right angles to the equipotential surface.*** To see why, suppose for a moment that it were not so, and that the field at some point on a surface had a part of itself lying flat along the surface. That flat part would be a force, and it would push a charge resting there sideways, along the surface. Let the charge move a little way in that direction, and the field, pushing as the charge moves, would do work upon it. But the charge has stayed on the equipotential, where the potential does not change, and the work done on a charge is the charge multiplied by the change of potential it passes through. With no change of potential there can be no work. The field cannot both do work and do no work on the same small journey. The only escape from the contradiction is to deny the thing that was supposed: the field has no part lying along the surface at all. Whatever the field does, it does straight out of the surface, exactly perpendicular to it. On the diagrams this shows as field lines piercing the equipotentials cleanly at ninety degrees, never grazing along them.

The second property is the one just used to prove the first, now promoted to a headline in its own right: ***no work is done in moving a charge from one point of an equipotential surface to another point of the same surface***. The reasoning is a single line. The work done by an outside agent in carrying a charge between two points is the charge multiplied by the difference of potential between those points, the relationship $W = q\Delta V$. Both points lie on the one equipotential, so they are at the same potential; the difference of potential between them is therefore zero, and the work, being the charge times that difference, is zero along with it. It does not matter how long the path is, how sharply it twists, or whether the charge is large or small, positive or negative. If the journey begins and ends on the same equipotential, the work is zero. This is the electrical echo of the walker who returns to the same contour having spent no net effort against gravity.

The third property restores the message that was hidden in the spacing of the map. ***Where equipotential surfaces are crowded close together the field is strong, and where they are spread far apart the field is weak.*** Imagine two neighbouring equipotentials whose potentials differ by some fixed step, say ten volts. If the two surfaces lie close together, those ten volts are used up over a very short distance, so the potential is changing quickly as one moves through the region. If the surfaces lie far apart, the same ten volts are spread thinly over a long distance, and the potential changes only slowly. As the next part of this section will shortly make exact, a potential that changes quickly with distance means a strong field, and a potential that changes slowly means a weak one. Crowded equipotentials, strong field; widely spaced equipotentials, weak field. The map told the truth about hills, and it tells the same truth about fields.

The fourth property is a short one. ***Two different equipotential surfaces can never cross one another.*** The argument is the very one that forbade two field lines from crossing, back in the discussion of the point charge. If two equipotentials of different value did meet at some point, that single point would lie on both surfaces at once, and so would be obliged to hold both values of potential at the same time. But the potential at a point in space is one single definite number; a place cannot be at twelve volts and at twenty volts together. The crossing is impossible, and the equipotentials, like well-mannered contour lines, thread their way among one another without ever touching.

The fifth property ties the whole idea back to conductors. ***The surface of a conductor sitting in electrostatic equilibrium is an equipotential surface,*** and more than that, the entire solid body of the conductor is at one and the same potential. The

reason was half built already in the work on Gauss's law. Inside a conductor in equilibrium the electric field is zero, for if any field remained the free charges would still feel a force and would still be drifting, and the conductor would not yet be in equilibrium at all. With the field zero everywhere inside, no work is needed to carry a charge from any one interior point to any other, and where no work is needed to move between two points those two points must be at the same potential. Every point of the metal, the surface and the depths alike, therefore shares one common value. A conductor in equilibrium is an equipotential body wrapped in an equipotential skin, and this is precisely why the field lines met its surface at right angles in the earlier figures: that surface is an equipotential, and the first property demands it.

So far the equipotential surfaces have been read like a map, by eye. It is time to make the reading exact, and to do that the section returns once more to the steepness of the hill. The steepness at any spot is the answer to a precise question: *as I take one short step across the contours, by how much does my height change?* A small change of height for the step means gentle ground; a large change means a cliff. The electric field, it now turns out, is the answer to exactly the same question asked of the potential: *as I take one short step through the field, by how much does the potential change?* The field is the steepness of the potential landscape, and the time has come to write that down as an equation.

In the previous section the potential was built out of the field by adding up the effect of the field all along a path, an operation of accumulation, an integral. The reverse operation, travelling from the potential back to the field, must therefore be an operation of rate of change. For a field that points along a single direction, that relationship is:

$$\mathbf{E} = -\frac{\mathbf{dV}}{\mathbf{dr}}$$

The quantity $\frac{\mathrm{dV}}{\mathrm{dr}}$ in that equation is the rate at which the potential changes with distance, and it carries a name of its own, the **potential gradient**. The equation says, in a few symbols, everything the map said in pictures: *the electric field at a point is equal to the negative of the potential gradient there.*

Two parts of that statement repay a careful look, the size and the sign. Take the size first. The equation says the field is large wherever the potential is changing rapidly with distance, and small wherever the potential changes only sluggishly. That is the third property over again, no longer dressed as a picture of crowded lines but written out as a piece of arithmetic: closely spaced equipotentials simply mean a large potential gradient, and a large potential gradient means a strong field. The strength of the field and the steepness of the potential are not two separate facts but one single fact seen twice.

Now the sign. The minus sign standing in front is not decoration; it fixes the direction in which the field points. It declares that the field points the way the potential decreases, from high potential towards low, downhill on the landscape of potential. Set a positive charge free in a field and it rolls, exactly as a ball set down on a hillside rolls, from high potential to low. The field is the arrow that shows which way is downhill, and the minus sign is the small but essential piece of bookkeeping that keeps that arrow pointing the correct way.

A check on units is worth a moment, because it heals an apparent split in the language of the subject. The potential gradient is a potential divided by a distance, so its unit is the volt per metre, written V/m. Yet the electric field was first introduced, chapters and sections ago, as a force per unit charge, measured in newtons per coulomb, written N/C. The two units look unrelated, and are in fact identical. A volt is a joule per coulomb, and a joule is a newton metre, so a volt per metre is a newton metre per coulomb, divided again by a metre, and the metre cancels to leave simply a newton per coulomb. One volt per metre is one newton per coulomb, $1\mathrm{V/m} = 1\mathrm{N/C}$. The force picture and the energy picture have been describing the very same field, in two accents, all along.

The relationship runs in both directions. Just as the field can be found from the potential by taking a gradient, the potential can be recovered from the field by the accumulation met in the previous section,

$$\mathrm{V_B} - \mathrm{V_A} = -\int \mathrm{E} \cdot \mathrm{dr},$$

the field summed all along a path from one point to the other. The words all along a path are the important ones, and they were earned in the previous section: because the electrostatic force is conservative, the sum comes out the same whatever route is chosen between the two points. The two equations, $\mathrm{E} = -\frac{\mathrm{dV}}{\mathrm{dr}}$ and its inverse, form a matched pair, each one undoing the work of the other. Together they hand the student a freedom well worth having: a problem may be attacked from the field end or from the potential end, whichever proves the easier, and the pair of equations then carries the answer safely across to the other description.

The simplest case of all is the uniform field between parallel plates, and it is worth seeing how the grand relationship collapses neatly onto the modest formula already in hand. Between the plates the potential does not curve; it changes at a perfectly steady rate from the one plate to the other. A steady rate of change means the potential gradient is just the whole potential difference divided by the whole separation, and the field is then the same everywhere, with magnitude:

$$\mathrm{E} = \frac{\mathrm{V}}{\mathrm{d}}$$

That small equation, first met when parallel plates were introduced, now stands revealed for what it truly is: not a separate rule to be memorised on its own, but the gentlest possible example of $E = -\frac{dV}{dr}$, a straight-line potential giving a constant gradient giving a constant field.

One question, more than any other, decides whether a student has truly grasped the relationship between field and potential, and it can be put in two halves.

First: *if the electric field is zero everywhere throughout some region of space, must the potential be zero there too?*

Second, the mirror of the first: *if the potential is zero at some point, must the field be zero there?*

The honest answer to both halves is no, and the contour map explains the no before a single charge has been drawn.

Think once more of the hillside. A wide flat plateau has zero slope, the ground level across the whole of it, and yet that plateau may sit two thousand metres above the sea. Flatness is not lowness; a piece of ground can be perfectly level and still be very high. That settles the first half. And a particular spot may happen to lie exactly at sea level, its height reading zero, while the ground there plunges steeply down towards a cove; the height is zero but the slope is not. That settles the second half. Height and slope are simply different things, and either one of them may vanish while the other stands.

Translated into electricity, the potential is the height and the field is the slope. Suppose the field is zero throughout a region. The equation for the field then says the gradient of the potential is zero there, which means the potential does not change from place to place; it is constant. But constant is a long way from zero. The cleanest illustration is the hollow charged conductor met earlier under Gauss's law. Everywhere inside the hollow the field is exactly zero, and yet the potential inside is not zero at all: it holds, unchanged, the same value the charged surface itself has, carried into the interior precisely because there is no field in there to alter it. A region of zero field is a region of flat, constant potential, and that constant may be a thousand volts as easily as nothing.

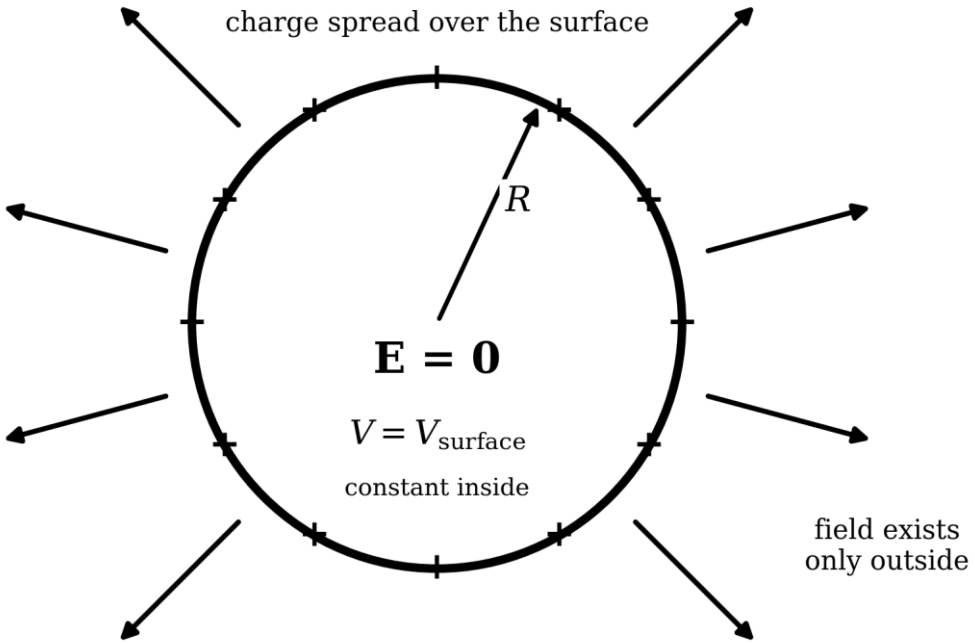


Figure: ***A hollow charged conductor.*** *The charge spreads itself over the outer surface. Inside the hollow the field is zero at every point, yet the potential is not zero: it holds the constant value of the charged surface all the way through the interior. Field lines exist only outside the conductor.*

The other half is settled by an object already familiar, the electric dipole of the previous section. At every point on the plane lying halfway between its two equal and opposite charges, the potential is zero, the positive contribution and the negative contribution being equal numbers of opposite sign that cancel exactly in the sum. Yet the field on that plane is not zero; it is strong and definite, because the field is a vector sum, and there the two field vectors of the two charges point so as to add together rather than to cancel. A point of zero potential is merely a place where signed numbers have happened to total zero; it makes no promise whatever about the field.

The lesson behind the two halves is one single lesson, and it is the very heart of this section. The field is obtained from the potential by differences and by rates of change, never from the bare value the potential happens to have at one point. The field equation has eyes only for how the potential changes; the actual figure the potential has reached is invisible to it. That is why a whole region sitting at one constant potential, however high that potential may be, contains no field, and why a lone zero of potential, with changing values pressing in all around it, can sit calmly in the middle of a powerful field.

Equipotential surfaces have now done their work. They have sorted the cluttered space around a charge into a clean set of level sheets, they have shown the field to be nothing grander than the steepness of those sheets, and they have kept the two ideas of field and potential firmly distinct even while binding them together with a single equation. Everything in the chapter so far, the field, the potential and the equipotential, has been a description of the empty space that charges drape around themselves. The next section at last lets something move. It places a charged particle into a uniform field and follows it as the field takes hold, accelerates it, bends its flight into a graceful curve, and, in one final and elegant trick, holds it perfectly still in the air against the steady pull of gravity.

Worked examples now, and the equipotential idea proves an unusually generous host: it has a habit of pressing an answer into the candidate's hand before the pen has properly warmed up. Four examples follow. Approach them in good spirits, for not a single one of them bites.

BINDER Example 28

A certain equipotential surface in an electric field has a potential of 12V at every point on it. A charge of +3.0C is carried from one point on this surface to another point on the same surface, by a slow and thoroughly winding route. Calculate the work done by the external agent in carrying the charge.

Solution

The work done in carrying a charge between two points is the charge multiplied by the difference of potential between those points:

$$W = q\Delta V$$

Both points lie on the same equipotential surface, so each of them is at a potential of 12V. The difference of potential between them is therefore:

$$\Delta V = 12V - 12V = 0V$$

Substituting this difference of potential, together with the charge, into the equation for the work:

$$W = q\Delta V = (3.0C)(0V) = 0J$$

The work done by the external agent is zero.

Making Sense of the Answer: *The long winding route was a deliberate distraction and changed nothing whatever. Work done on a charge depends only on the potentials of the starting point and the finishing point, never on the journey taken between them, and here those two potentials were equal. Had the charge instead been carried to a point on a different equipotential, a surface at twenty volts say, the answer would not have been zero. It is the sameness of the two potentials, and not the motion itself, that drives the work down to nothing.*

Think Like a Physicist: *When a question places the start and the finish of a journey on the same equipotential surface, write zero for the work at once, and read the rest of the question only out of politeness. The length of the path, the number of its bends and the size of the charge are all decoration laid over a single fact: equal potentials, no work.*

BINDER Example 29

Two large parallel metal plates are mounted horizontally, one above the other, separated by a gap of 4.0cm. The upper plate is held at a potential of 200V and the lower plate at 0V.

(a) Find the magnitude of the uniform electric field in the gap between the plates.
(b) State the potential of the equipotential surface lying exactly halfway between the plates, and find how far above the lower plate the 50V equipotential surface lies.

Solution

(a) Between parallel plates the field is uniform, which means the potential changes at a perfectly steady rate across the gap and the potential gradient is simply the whole potential difference divided by the whole separation:

$$E = \frac{\Delta V}{d}$$

The potential difference between the plates is 200V − 0V = 200V, and the separation is 4.0cm, which is 0.040m. Substituting these values:

$$E = \frac{200V}{0.040m} = 5.0 \times 10^3 V/m$$

The field between the plates has magnitude 5.0×10^3V/m, which is the same thing as 5.0×10^3N/C.

(b) Because the field is uniform, the potential rises in step with distance: equal climbs of distance bring equal climbs of potential. Halfway across the gap the potential has therefore reached exactly half of its full range, so the equipotential surface there carries a potential of:

$$\frac{200V + 0V}{2} = 100V$$

To locate the 50V surface, measure a height upward from the lower plate. The potential at that height, rising steadily from zero at the lower plate at a rate equal to the field strength, is the field strength multiplied by the height, $V = Ex$, so the height is the potential divided by the field strength:

$$x = \frac{V}{E} = \frac{50V}{5.0 \times 10^3 V/m} = 0.010m$$

The 50V equipotential surface lies 0.010m, that is to say 1.0cm, above the lower plate.

Making Sense of the Answer: *The equipotentials between the plates are evenly spaced, and that even spacing is the whole signature of a uniform field. The 50V surface sits a quarter of the way up the gap for the plainest of reasons: 50V is one quarter of the 200V total, and in a uniform field equal steps of potential always occupy equal steps of distance. The numbers have simply shared the gap out in proportion.*

Think Like a Physicist: *In a uniform field the potential behaves like a ruler laid flat across the gap, its markings spaced evenly from the one plate to the other. Locating a particular equipotential is then not really a piece of physics at all but a piece of proportion: find what fraction the wanted potential is of the total, and the surface sits that very same fraction of the way across the gap.*

REAL Example 30

During a storm a high-voltage TANESCO line is brought down, and one end of it comes to rest on the open ground of a field. The point where the wire touches the earth is now held at a high potential, and the potential of the ground falls away steadily with distance from that point, so that the earth around the fallen wire is covered with a set of equipotential surfaces shaped like rings centred on the point of contact. A person caught in that field is endangered not by touching the wire but by the ground itself.

(a) Using the idea of a potential gradient, explain why a person who walks away with long, normal strides is in greater danger than one who shuffles away with the feet kept close together.
(b) Explain why, once the danger is understood, the safest way to cross such ground is to move with the two feet always touching or kept together.

Solution

(a) What threatens a person standing on that ground is not the potential of the ground itself but the difference of potential between the two feet, for it is that difference, acting as a voltage applied across the body, that drives a current up one leg and down the other. The potential of the ground changes with distance from the point of contact, so two feet planted at two different distances stand on two different equipotential rings and are held at two different potentials. A long stride sets the feet far apart. The potential gradient, the rate at which the potential changes with distance, then acts across that long gap between the feet, and a gradient multiplied by a long distance gives a large difference of potential, and so a large and dangerous current through the body. A short shuffling step sets the feet almost together. The same gradient now acts across only a tiny distance, the two feet sit on very nearly the same ring at very nearly the same potential, the difference between them is small, and only a small and survivable current can flow. Cutting the distance between the feet cuts the voltage across the body in the very same proportion.

(b) The principle behind the danger points straight at the remedy. Since the hazard is the difference of potential between the feet, and since that difference grows with the distance between the feet, safety lies in never letting the feet be far apart. Ordinary walking is dangerous because at every stride the feet are spread wide across many rings, and running is worse still, for it throws the feet far apart and lands them hard on widely separated rings. The safe method is the exact opposite of a stride: keep the two feet touching one another and move by sliding them forward together in small shuffles, or by hopping with the feet kept firmly together, so that at every instant both feet rest on practically the same equipotential ring and practically no difference of potential is ever allowed to appear across the body.

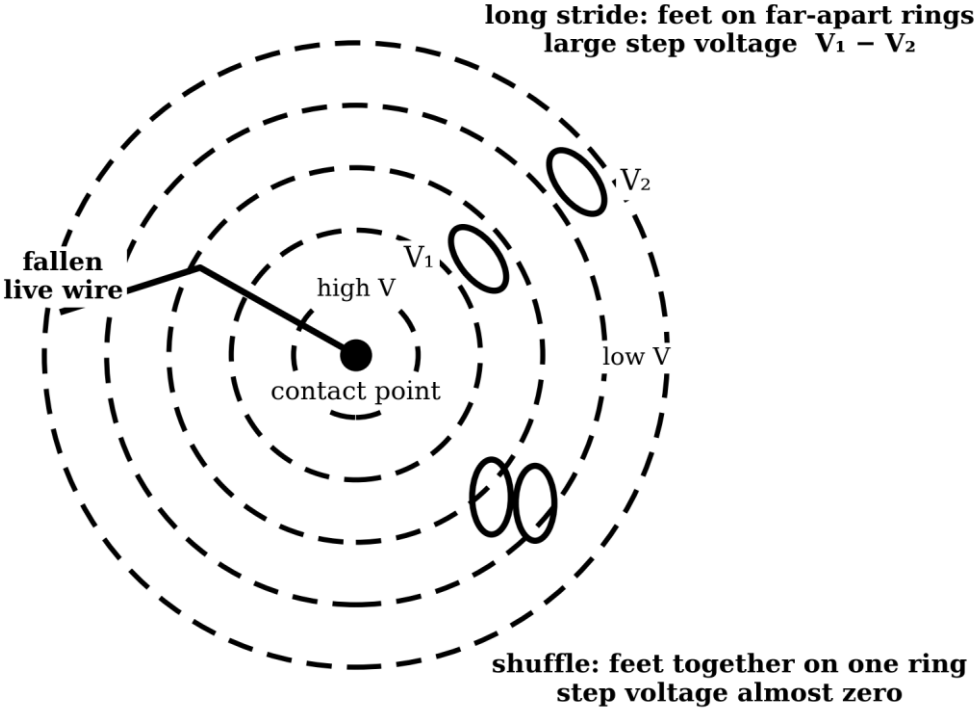


Figure: ***Step potential near a fallen conductor.*** *Seen from above, the ground around the point of contact is covered with equipotential rings. Two feet placed a long stride apart rest on widely separated rings, and a large step voltage appears between them. Feet kept together rest on a single ring, and almost no step voltage appears.*

Making Sense of the Answer: *The danger is carried by the difference of potential, and not by the potential itself, the very same lesson taught by the bird that perches unharmed on a single high-voltage line. Here that lesson is turned into a rule for survival. The equipotential rings lying on the ground cannot be seen, but they are every bit as real as the contour lines on a map, and a person's safety is decided entirely by how many of those rings lie trapped between the two feet.*

Think Like a Physicist: *A hazard that announces itself as a high voltage is very often, when it is looked at properly, a hazard about a voltage difference taken across a small span. Reduce the span and the difference shrinks along with it. Reading the danger through the idea of a potential gradient turns a frightening and confusing situation into a simple matter of geometry: keep the two sampling points, here the two feet, close together, and the gradient is left with only a short distance over which to build a difference.*

HOT Example 31

Two students are arguing.

The first says: *If the electric field is zero everywhere inside some region of space, then the electric potential must be zero everywhere in that region as well.*

The second disagrees: *A zero field only means the potential is the same at every point of the region, and that same value may be anything at all, a large value included.*

To settle the matter they consider a hollow metal sphere of radius 0.10m carrying a charge of +8.0nC spread over its surface. Decide which student is correct, and support the decision by finding both the electric field and the electric potential at the centre of the sphere.

Solution

Two separate facts about a charged hollow conducting sphere are needed, and both were established earlier in this chapter with the help of Gauss's law.

Take the field first. A Gaussian surface drawn anywhere inside the hollow sphere encloses no charge whatever, since all of the charge sits on the outer surface of the metal. Gauss's law then gives zero for the field at every interior point, the centre included:

$$E_{centre} = 0$$

Now take the potential. Outside the sphere the potential is that of a point charge, $V = \frac{kQ}{r}$.

Following it inward, it rises steadily as the distance shrinks until, at the surface, where the distance equals the radius of the sphere, it reaches the value $\frac{kQ}{R}$.

Inside the surface the field is zero, so there is no gradient anywhere in the interior to change the potential any further; it therefore holds that same surface value, unchanged, all the way to the centre:

$$V_{centre} = \frac{kQ}{R}$$

Substituting:

$$V_{centre} = \frac{(9.0 \times 10^9 Nm^2/C^2)(8.0 \times 10^{-9} C)}{0.10m} = 7.2 \times 10^2 V$$

At the centre of the sphere the field is zero while the potential is 720V. The field has vanished; the potential plainly has not. The first student is wrong and the second student is right.

Making Sense of the Answer: *The result is the hollow-conductor fact and the dipole fact pressed together into one single lesson. A field of zero does not empty the potential out; it only flattens it, fixing it at one constant value throughout the region. Field and potential are independent enough that either one of them may stand at zero while the other does not, because the field measures the gradient of the potential, and a gradient of zero is the mark of a constant, not the mark of a nothing.*

Think Like a Physicist: *When a question sets a field against a potential, refuse to let the one quietly stand in for the other. Ask two separate questions. For the field, the only question is whether the potential is changing at the point; if it is not, the field is zero there. For the value of the potential, the question is what the potential has added up to at that point, which is a quite separate matter. The hollow charged sphere is the cleanest place in all of electrostatics to keep those two questions apart, and a student who can keep them apart there can keep them apart anywhere.*

Four examples, and not a single casualty among them. Equipotential surfaces, for all the quiet labour they save a person, turn out to be about the most peaceable and best-mannered objects in the whole of electrostatics.

MOTION OF A CHARGED PARTICLE IN A UNIFORM ELECTRIC FIELD

Every page of this chapter so far has drawn the electric field as a kind of landscape, a fixed pattern of arrows and level surfaces hanging in empty space, waiting. Nothing in it has actually moved. That changes now. This section takes a single charged particle, sets it down inside a uniform field, and simply watches what the field does to it. And the most useful thing to carry into that watching is a comparison the reader already owns. A uniform electric field does to a charge exactly what the uniform gravity near the ground does to a thrown stone. Every result about projectiles, every parabola traced by a thrown ball, every trick of splitting a motion into a horizontal part and a vertical part, carries straight across. The physics of this section is, in a real sense, already familiar to the reader; only the names have changed.

The natural place to stage that comparison is the region between two parallel charged plates, because there, and almost nowhere else so cleanly, the field is uniform. The field E has the same size and the same direction at every point of the gap, so a charge placed anywhere in that gap meets one and the same field. As the next few paragraphs show, that single fact of uniformity is what makes the motion so simple, and so very like the flight of a stone.

Begin with the force. A charge Q sitting in a field E feels an electric force, and the size of that force was settled in the very first pages of the chapter, in the definition of the field itself:

$$F = QE$$

In a uniform field E is constant, so this force is constant too, the same in size and the same in direction wherever the charge happens to lie in the gap. A constant force is the key that unlocks everything else in the section.

For a constant force acting on a particle of mass m, Newton's second law gives a constant acceleration. Dividing the force by the mass:

$$\mathbf{a = \frac{F}{m} = \frac{QE}{m}}$$

The acceleration of a charged particle in a **uniform** field is **constant**. That one sentence is the whole engine of the section. A constant acceleration is precisely the condition under which all the equations of uniformly accelerated motion, the ones already learned for objects falling under gravity, apply without the smallest change.

The charge in the field is, mathematically, a projectile, and the quantity $\frac{QE}{m}$ simply takes the part that g took for falling bodies.

Two remarks before the motion is set going.

First, the direction. A positive charge is pushed along the field and accelerates in the direction of E; a negative charge, an electron above all, is pushed against the field and accelerates in the direction opposite to E.

Second, a point of bookkeeping. The particle has mass, so the Earth pulls on it with a weight mg as well. Whether that weight matters depends entirely on what the particle is. For an electron or an ion the mass is so fantastically small that the weight is utterly swamped by the electric force, thousands upon thousands of times smaller, and it is dropped without apology.

For a larger object, a charged drop of oil say, the weight is comparable with the electric force and cannot be ignored at all. This section meets both cases, and the difference between them is simply the difference between an electron and an oil drop.

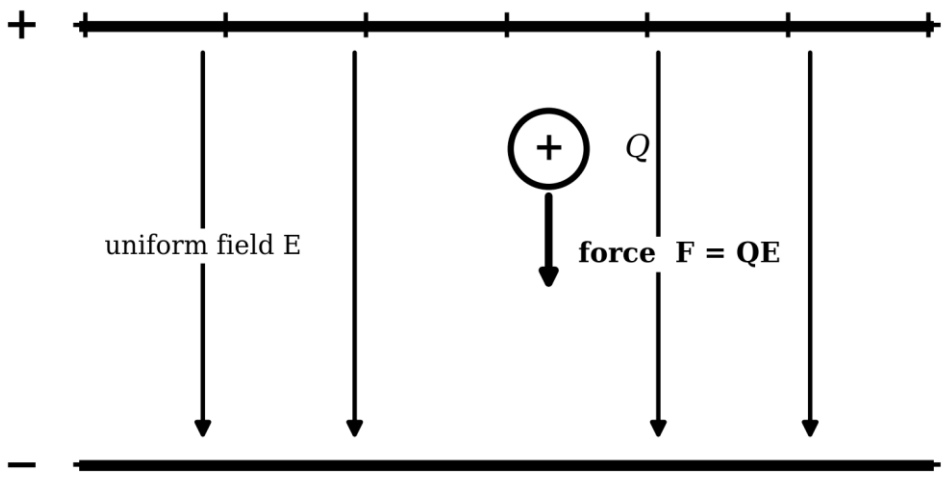


Figure: ***A charge in a uniform field.*** *Between parallel plates the field is uniform, so the electric force on a charge has the same size and direction everywhere in the gap. A constant force produces a constant acceleration, exactly as a constant gravitational field does for a falling body.*

The simplest motion is the one in which the charge is released from rest, or is already travelling straight along a field line. With the force directly along the motion, the charge runs in a straight line and simply speeds up, gaining velocity at the steady rate $a = \frac{QE}{m}$, exactly as a stone dropped from the hand gathers speed at the rate g. There is nothing fresh to derive: the familiar equations of straight-line accelerated motion deliver the velocity gained and the distance covered. The energy side of it was settled already in the work on potential, where the push of the field across a potential difference was shown to hand the charge a kinetic energy equal to Q times the potential difference crossed. Speeding a charge up in a straight line is the principle of the electron gun that launches the beam in every cathode-ray tube. It is motion along the field, and it holds no surprises.

The interesting motion is the one across the field. Let a charged particle be fired into the uniform field so that it enters moving at right angles to the field, travelling parallel to the plates rather than along the gap between them. This is the exact counterpart of a stone thrown horizontally from the top of a cliff, and it is handled in the exact same way: split the motion into two independent parts and treat each part entirely on its own.

Set an axis x along the direction the particle was fired, and an axis y along the field. Along x there is no force at all, because the field, and so the force QE, points entirely along y. With no force along x, the particle simply keeps whatever speed u it entered with, and after a time t it has travelled a distance:

$$x = ut$$

Along y the story is the story of the falling stone. The particle starts with no velocity in the y direction, and the constant force gives it the constant acceleration $a = \frac{QE}{m}$. Starting from rest and accelerating steadily for the same time t, it is carried sideways a distance:

$$y = \frac{1}{2}at^2$$

Two motions, but one clock. The time t is the single thread running through both, and eliminating it joins the two motions into the shape of the path.

From the first equation the time is $t = \frac{x}{u}$. Substituting that into the second equation, and then replacing the acceleration by $\frac{QE}{m}$, gives

$$\mathbf{y = \frac{1}{2}a\left(\frac{x}{u}\right)^2 = \frac{QEx^2}{2mu^2}}$$

The crucial thing in that result is not the cluster of constants but the shape. The deflection y is proportional to x^2, and a relationship in which one coordinate grows as the square of the other is a parabola. The charged particle, crossing the uniform field, curves along a parabola, just as surely as the stone thrown from the cliff traces one under gravity.

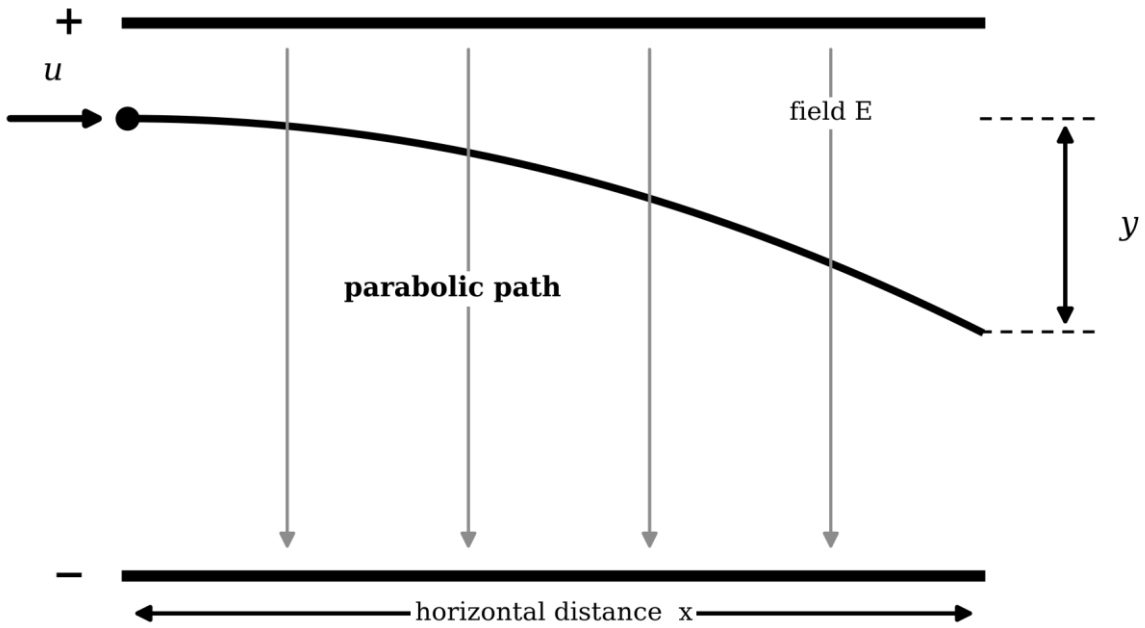


Figure: ***The parabolic path.*** *A particle fired into the field at right angles to it keeps a steady speed along its original direction while being steadily accelerated along the field. The two motions combine into a parabola, the same curve a stone traces when thrown horizontally.*

This parabola is not a curiosity. It is the working principle of a whole family of instruments. Steer a beam of electrons between a pair of charged plates and the size of the deflection can be read straight off the parabola; vary the voltage on the plates and the beam can be aimed at any chosen spot on a screen. That is the cathode-ray oscilloscope tracing its glowing line, the older television painting its picture, and the inkjet printer flicking charged droplets of ink to their appointed places on the page. Every one of them is this paragraph's parabola, put to work.

The last motion is the most surprising of all, because in it nothing moves. Take the charged oil drop mentioned earlier, the case in which the weight is not negligible, and set it between two horizontal plates with the field running vertically. Two forces now act on the drop: Its **weight mg** pulling it straight down, and the **electric force QE** which, if the field is pointed correctly, pushes it straight up.

Adjust the strength of the field, by adjusting the voltage on the plates, until the upward electric force exactly matches the downward weight. At that one setting the two forces cancel, the total force on the drop is zero, and the drop hangs in the air, perfectly still, suspended on nothing but the balance of electricity against gravity.

The condition for that **motionless balance** is simply that the two forces are equal:

$$\mathbf{QE = mg}$$

It is a short equation, but one of the most quietly powerful in all of physics, because of what can be done with it. Every quantity in it can be measured. The weight side, mg, can be found by other means; the field E is known from the plate voltage and the plate spacing. That leaves the charge Q as the only unknown, and the equation hands it straight over.

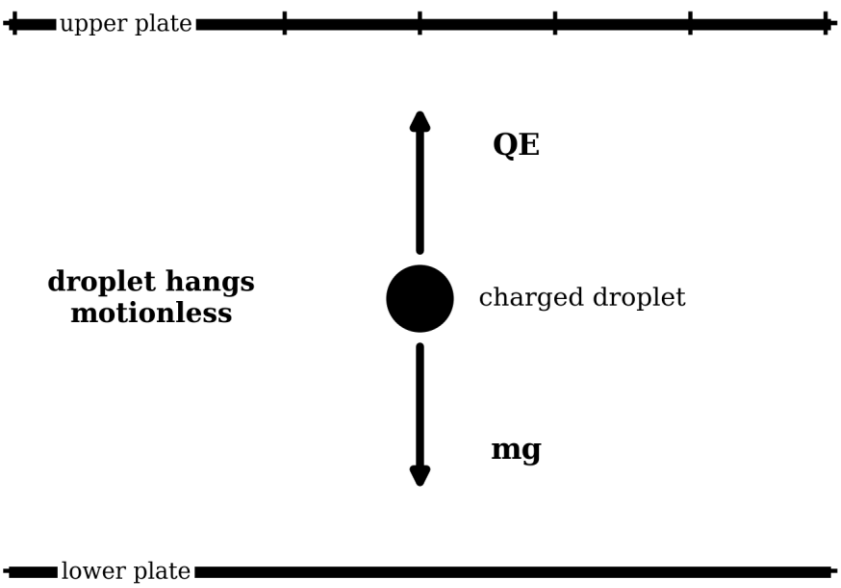


Figure: ***The balancing act.*** *A charged droplet between horizontal plates feels its weight downward and the electric force upward. When the field is adjusted so the two forces are equal the drop hangs motionless, and the balance condition becomes a way of measuring the charge.*

This is precisely the experiment Robert Millikan performed (**Millikan oil-drop experiment**), beginning in 1909, with tiny drops of oil. He sprayed a fine mist of oil between two horizontal plates, picked out a single charged droplet under a microscope, and adjusted the voltage until that droplet hung motionless. From $\mathrm{QE} = \mathrm{mg}$ he then calculated the charge on the drop. Repeating the measurement on droplet after droplet, he found something remarkable: the charge was never just any value, but always a whole-number multiple of one smallest basic amount. That basic amount was the charge on a single electron, the elementary charge e, the very quantisation of charge met at the start of this chapter.

Four worked examples follow, and they share a single secret: every one of them is a projectile problem wearing a disguise. A reader who once survived the mechanics of thrown stones already has, tucked away, everything needed to finish all four. The disguises are rather good, but they are only disguises.

BINDER Example 32

A uniform electric field of strength $5.0 \times 10^{3}\mathrm{N/C}$ is set up between two parallel plates. A small charged particle, of charge $2.0 \times 10^{-6}\mathrm{C}$ and mass $4.0 \times 10^{-8}\mathrm{kg}$, is placed in the field.

(a) Find the electric force acting on the particle.

(b) Find the acceleration this force produces.

Solution

(a) The electric force on a charge in a field is the charge multiplied by the field strength:

$$\mathrm{F} = \mathrm{QE}$$

Substituting the charge and the field strength:

$$\mathrm{F} = \mathrm{QE} = (2.0 \times 10^{-6}\mathrm{C})(5.0 \times 10^{3}\mathrm{N/C}) = 1.0 \times 10^{-2}\mathrm{N}$$

(b) The acceleration follows from Newton's second law, the force divided by the mass:

$$\mathrm{a} = \frac{\mathrm{F}}{\mathrm{m}} = \frac{1.0 \times 10^{-2}\mathrm{N}}{4.0 \times 10^{-8}\mathrm{kg}} = 2.5 \times 10^{5}\mathrm{m/s^2}$$

The particle accelerates at $2.5 \times 10^{5}\mathrm{m/s^2}$ in the direction of the field.

Making Sense of the Answer: *The acceleration is enormous, hundreds of thousands of metres per second in every second, and yet the force that produces it, a hundredth of a newton, would not trouble a housefly. It is the smallness of the mass that magnifies a gentle force into a violent acceleration. It is worth noticing, too, that the weight of this particle, its mass times g, is only about $4 \times 10^{-7}N$, some twenty-five thousand times smaller than the electric force. For a particle this light, gravity is rightly left out of the account.*

Think Like a Physicist: *A problem of this kind is two steps, never one. The first step, $F = QE$, belongs to electricity; the second step, $a = \frac{F}{m}$, belongs to mechanics and has nothing electrical in it at all. Keeping the two steps mentally apart, the*

field producing a force and then the force producing an acceleration, is what keeps a charged-particle problem from ever turning into a muddle.

REAL Example 33

In an experiment modelled on Millikan's, a tiny charged oil drop of mass $3.3 \times 10^{-15}\text{kg}$ is held perfectly motionless in the uniform field between two horizontal plates. The field has strength $2.0 \times 10^{5}\text{N/C}$ and is directed so that the electric force on the drop is upward.

(a) State the condition for the drop to hang motionless, and use it to find the charge on the drop.

(b) Taking the elementary charge as $1.6 \times 10^{-19}\text{C}$, find how many excess electrons the drop carries.

Solution

(a) For the drop to hang motionless the total force on it must be zero, which means the upward electric force must exactly balance the downward weight:

$$\text{QE} = \text{mg}$$

Making the charge the subject and substituting the values, with g taken as 9.8m/s^2:

$$\text{Q} = \frac{\text{mg}}{\text{E}} = \frac{(3.3 \times 10^{-15}\text{kg})(9.8\text{m/s}^2)}{2.0 \times 10^{5}\text{N/C}} = 1.6 \times 10^{-19}\text{C}$$

(b) The number of excess electrons is the charge on the drop divided by the charge on a single electron:

$$\text{n} = \frac{\text{Q}}{\text{e}} = \frac{1.6 \times 10^{-19}\text{C}}{1.6 \times 10^{-19}\text{C}} = 1$$

The drop carries a single excess electron.

Making Sense of the Answer: *The charge came out equal, to the precision of the measurement, to the elementary charge itself, so the drop carries exactly one electron more than its protons can balance. This is the very heart of Millikan's discovery. Drop after drop returned a charge that was always a whole number of these elementary units and never a fraction of one, and that stubborn whole-number pattern was the proof that electric charge comes in indivisible grains.*

Think Like a Physicist: *The balance condition turns a stationary object into a measuring instrument. Nothing moves, and yet because the drop is motionless a physicist may write down at once that the forces upon it are equal, and an equation with a single unknown is an equation already half solved. Whenever an object is described as hanging, floating, or suspended, the first line of the answer should record that the forces upon it sum to zero.*

REAL Example 34

Kipanga and Kipute are watching a demonstration in which charged particles are sent into a uniform electric field. Kipanga argues that a positive ion and an electron, sent into the very same field, must come away with accelerations of equal magnitude, since the field is identical for both and the field is what does the pushing. Kipute insists that the two accelerations are nothing alike. Decide which student is right, and explain the reasoning.

Solution

Both students have hold of a piece of the truth, which is what makes the disagreement worth settling with care. Begin with the force. The electric force on a particle is QE. An electron and a singly charged positive ion carry the same magnitude of charge, the elementary charge e, and they are placed in the same field E, so the force on each of them genuinely does have the same magnitude. To that extent Kipanga is right, and it is presumably this equal force that the argument has in mind. But a force is not an acceleration. The acceleration is the force divided by the mass, $\text{a} = \frac{\text{QE}}{\text{m}}$, and here the two particles part company entirely. An electron is among the lightest objects in all of nature; even the lightest possible ion, a lone proton, is already about 1.8×10^{3} times more massive, and a typical ion thousands of times more massive still. Dividing the same force by these wildly different masses produces wildly different accelerations. Set side by side, the ratio of the ion's acceleration to the electron's is:

$$\frac{\text{a}_{\text{ion}}}{\text{a}_{\text{e}}} = \frac{\text{m}_{\text{e}}}{\text{m}_{\text{ion}}},$$

the inverse ratio of their masses, which is a tiny fraction. The electron, for the same force, is flung forward with an acceleration tens of thousands of times greater than the ion's. Kipute is right: equal force, but utterly unequal acceleration. Kipanga's slip was to let the word force quietly turn into the word acceleration somewhere in the middle of the argument.

Making Sense of the Answer: *The resolution is the equation* $a = \frac{QE}{m}$ *read with care. The field fixes the force; the mass then fixes how much motion that force can produce. Two particles in one field feel matched forces and answer with mismatched accelerations, in exact inverse proportion to their masses. This is the reason that, in every device which steers particles with fields, it is electrons that are steered: being the least massive, they are the most obedient to the field.*

Think Like a Physicist: *Whenever two objects are compared in the same field, separate the two questions the comparison hides inside itself. Is the force the same? That depends only on the charge and the field. Is the acceleration the same? That*

depends also on the mass. A great many confident wrong answers come from settling the first question and then forgetting that the second one was ever asked.

HOT Example 35

An electron is fired horizontally, at a speed of $3.0 \times 10^6 \text{m/s}$, into a uniform electric field of strength 200N/C directed vertically downward. The field fills a region 10cm long, measured along the electron's original direction of travel. Taking the electron's charge as $1.6 \times 10^{-19}\text{C}$ and its mass as $9.11 \times 10^{-31}\text{kg}$, find the vertical deflection of the electron by the time it leaves the field region.

Solution

The motion splits into two independent parts, one along the original direction of travel and one along the field.
Along the direction of travel there is no force, so the electron keeps its entry speed. So:

$$\text{t} = \frac{\text{x}}{\text{u}} = \frac{0.10\text{m}}{3.0 \times 10^6 \text{m/s}} = 3.33 \times 10^{-8}\text{s}$$

Along the field the electron starts from rest and accelerates. The acceleration is the electric force divided by the mass:

$$\text{a} = \frac{\text{QE}}{\text{m}} = \frac{(1.6 \times 10^{-19}\text{C})(200\text{N/C})}{9.11 \times 10^{-31}\text{kg}} = 3.51 \times 10^{13}\text{m/s}^2$$

The vertical deflection is the distance the electron is carried, from rest, under this acceleration in the time just found:

$$\text{y} = \frac{1}{2}\text{at}^2 = \frac{1}{2}(3.51 \times 10^{13}\text{m/s}^2)(3.33 \times 10^{-8}\text{s})^2 = 1.95 \times 10^{-2}\text{m}$$

The electron leaves the field deflected by $1.95 \times 10^{-2}\text{m}$, which is very nearly 2.0cm.

Making Sense of the Answer: *The deflection is small, around two centimetres, and yet it is built from an acceleration of staggering size, $3.51 \times 10^{13} m/s^2$. What keeps the deflection so modest is the brevity of the crossing: the electron is inside the field for only about thirty nanoseconds, and even a colossal acceleration can move a particle only so far in so short a time. The same answer follows in a single line from the parabola formula $y = \frac{QEx^2}{2mu^2}$, which is just this two-step calculation compressed into one.*

Think Like a Physicist: *The whole method rests on refusing to think about the curved path directly. The parabola looks complicated; its two shadows, one cast along the direction of travel and one along the field, are each as simple as anything in mechanics, a steady speed and a steady acceleration from rest. Solve the two simple shadows on their own, let the shared time link them at the end, and the complicated curve has been conquered without ever being faced head on.*

With that, the moving charge has been followed through its three motions: straight down the field with a steady acceleration, across the field along a parabola, and held in a dead balance against gravity. The field has been mapped, the potential has been measured, the equipotentials have been drawn, and now a charge has been set loose in the field and watched. The theory of the electrostatic field is, with this section, complete. What remains is to see it earn its living.

The next section leaves the tidy world of plates and test charges and walks out into the ordinary one, into the machines and the hazards and the everyday wonders, from the lightning conductor on a roof to the photocopier humming in an office, that are nothing more than the electrostatics of this chapter quietly at work.

APPLICATIONS OF ELECTROSTATICS

Every idea in this chapter so far has been built and tested in the clean, quiet world of the textbook, with its point charges and its perfect plates. It is time to open the door and let the chapter out into the ordinary world, where it turns out to have been working all along, unremarked, inside machines and weather and the dust of a cement works. Seven applications follow, each a short look at one familiar thing through the eyes of this chapter. Not one of them needs an idea that has not already been met.

1. The Van de Graaff generator

This is the machine that makes a person's hair stand on end at a science fair, and it is electrostatics made visible. A moving rubber belt carries charge, sprayed onto it at the bottom, up into a large hollow metal sphere, where it is collected. Because the charge on a conductor always travels to the outer surface, met earlier in this chapter, the charge spreads over the outside of the sphere and the potential of the dome climbs higher and higher, reaching hundreds of thousands of volts.

The spraying of charge onto the belt relies on the point-discharge action, the same crowding of field at a sharp point that the lightning conductor will use below. A research version of exactly this machine, used as a particle accelerator, runs in physics laboratories all over the world. The principle, dome, belt, and sharp points, never changes.

2. The electrostatic precipitator

Inside the chimney of a cement works or a coal power station, a high-voltage wire runs down the centre of the gas duct. The strong field around the wire charges every speck of ash and smoke that drifts past it. Once charged, each particle feels the

force F = QE and is driven sideways onto earthed collecting plates, where it sticks and is later shaken off. Clean gas leaves the top of the chimney.

This is electrostatics doing a public service. The cement industry is a major one in Tanzania, and a works such as the Wazo Hill plant at Dar es Salaam would, without precipitators, blanket the surrounding district in fine grey dust. The whole device is nothing more than the force on a charge in a field, scaled up and put to work cleaning the air.

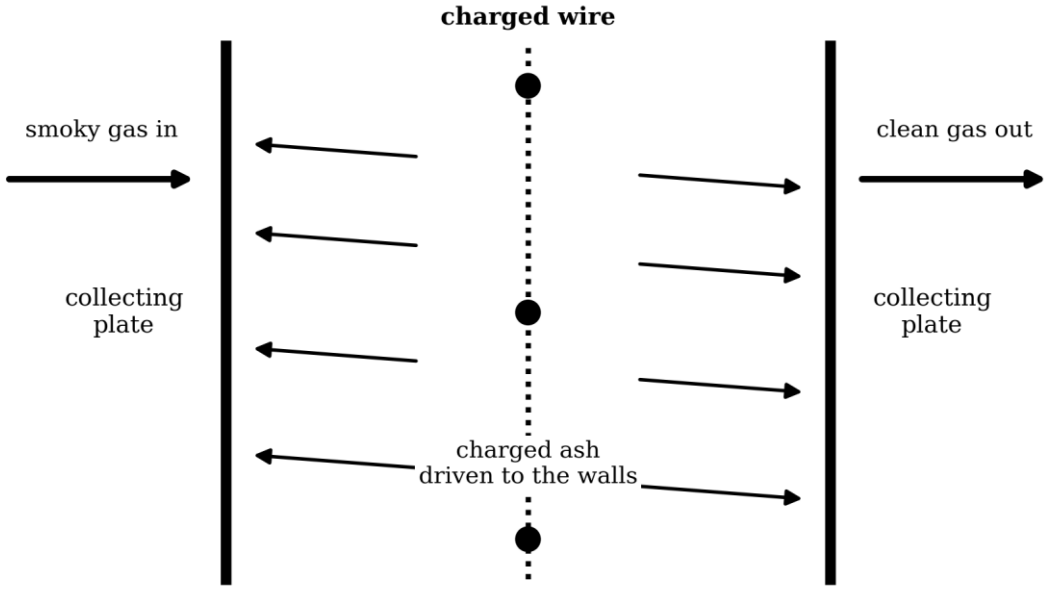


Figure: ***The electrostatic precipitator.*** *Smoky gas passes a charged wire that gives every ash particle a charge. The charged particles are then pulled sideways onto the collecting plates, and clean gas leaves the chimney.*

3. The photocopier and the laser printer

Both machines paint with charge. A metal drum is first given a uniform electrostatic charge in the dark. Light, reflected from the white parts of the page or written by a laser, lands on the drum and quietly removes the charge wherever it falls, leaving a hidden pattern of charge in the shape of the dark text and pictures alone.

Fine particles of toner, themselves charged, are then dusted over the drum. They cling by electrostatic attraction only to the charged regions, and so the invisible pattern becomes a visible one. The toner is pressed onto paper and fixed with heat. Every copy that slides out of the machine is a picture drawn entirely by the attraction and repulsion of charge.

4. The lightning conductor

A lightning conductor is a metal rod, sharply pointed at the top, fixed to the highest part of a building and joined by a thick copper strap to a metal plate buried in the earth. Its sharp point is the heart of it. Field crowds intensely at any sharp point on a conductor, a fact established earlier in the chapter, and the air there is ionised so strongly that charge leaks steadily and silently away into the sky, easing the build-up that would otherwise grow into a strike.

Should a strike come even so, the conductor offers it a broad, easy, metal path straight to the ground, and the dangerous current passes down the strap instead of through the brick and timber of the building. This is also why a tall coconut tree is struck in preference to a short mango tree beside it: greater height means a shorter air gap to the cloud, and a field that reaches the breakdown value of air sooner.

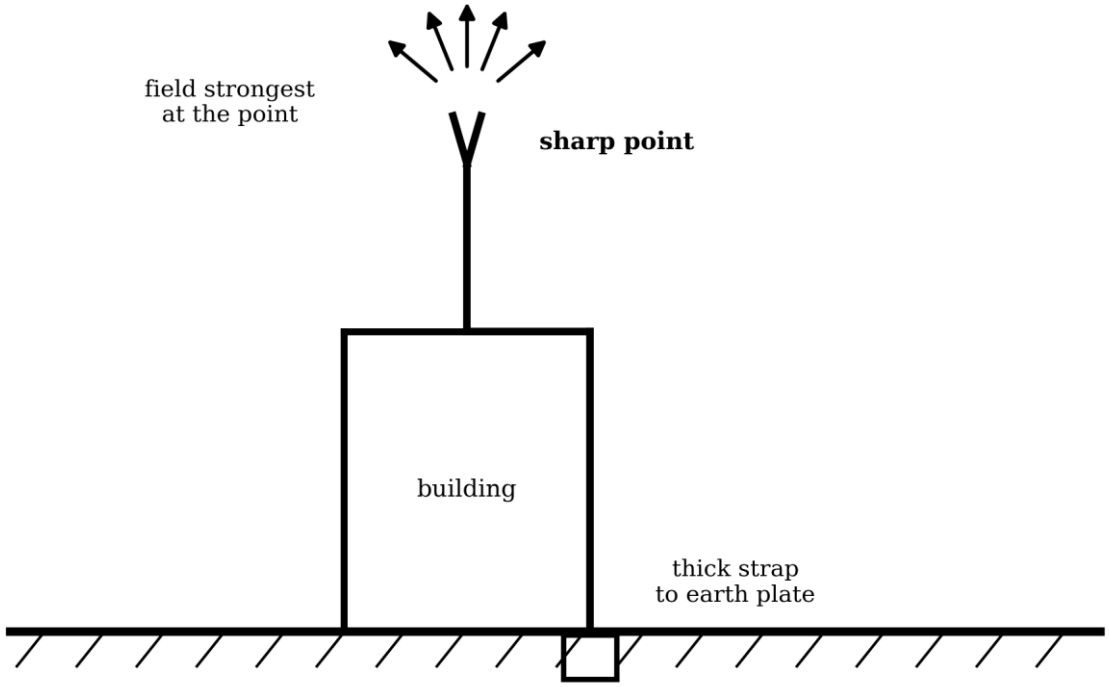


Figure: ***The lightning conductor.*** *A sharp point at the top of the conductor concentrates the electric field, so charge leaks quietly into the air there. A thick strap carries any strike safely down to an earth plate.*

5. The inkjet printer

An inkjet printer fires a fine stream of tiny ink droplets, and gives each droplet a controlled charge as it is formed. The charged droplets then fly between a pair of charged deflecting plates, where the uniform field bends each one along the parabola derived in the previous section. By setting the charge on a droplet, the machine chooses how far that droplet is deflected, and so steers it to its exact place on the page. It is the parabola of a charged particle in a uniform field, printing a letter.

6. Electrostatic spray painting

When a vehicle body is spray-painted in a factory, the paint leaving the spray gun is given an electric charge, and the metal body is earthed. Every charged droplet of paint is then attracted to the metal, and even droplets that would otherwise sail past are pulled back and curve around onto the far side of the panel. The result is an even coat, painted edges and hidden corners included, with very little paint wasted. The attraction between a charge and an earthed conductor, here, simply paints a car.

7. The static hazard in grain handling

Not every application is a benefit; the same physics also names a danger. On a farm or in a processing plant, dry grain pouring along a conveyor belt or down a chute rubs constantly against the surfaces it touches, and that friction separates charge exactly as a comb does through hair. Charge builds up on the insulated metal of the equipment, the potential climbs, and if it climbs far enough a spark jumps. In the dust-laden air above stored grain that spark can start a fire or an explosion.

The cure is the chapter in reverse: earth every piece of metal equipment with a conducting strap, so that charge drains harmlessly away as fast as friction deposits it and the potential is never allowed to build. Understanding electrostatics, here, is a matter of safety in an industry on which many Tanzanian livelihoods depend.

Seven machines and hazards, and a single chapter beneath all of them. The force on a charge, the crowding of field at a point, the parabola in a uniform field, the charge that flees to the surface of a conductor, the breakdown of air: each idea, built carefully in its turn, has been seen here earning its keep. The theory was never the destination. This was.

MISCELLANEOUS WORKED EXAMPLES ON ELECTROSTATICS

Example 36

(a) A glass rod that has been rubbed with silk is brought close to small pieces of dry uncharged paper. Explain why the pieces of paper are attracted to the rod even though they carry no net charge.

(b) Three small charged objects A, B and C are fixed at the vertices of an equilateral triangle of side 20cm. The charge at A is $+5.0\mu\text{C}$, at B is $-3.0\mu\text{C}$ and at C is $+4.0\mu\text{C}$. Calculate the magnitude and direction of the resultant electrostatic force on the object at A. Take $\text{k} = 9.0 \times 10^9 \text{N} \cdot \text{m}^2/\text{C}^2$.

Solution

(a) The positive rod polarises the neutral paper: electrons within the paper shift slightly toward the surface facing the rod, leaving the surface farther from the rod slightly positive. The two induced surfaces carry equal and opposite charge, but the negative one is closer to the rod than the positive one. The attraction between the rod and the nearer negative surface therefore exceeds the repulsion between the rod and the farther positive surface, and the net force on the paper is one of attraction.

(b) Place A at the origin, with B on the positive x-axis at $(0.20\text{m}, 0)$ and C above the x-axis at $(0.10\text{m}, 0.10\sqrt{3}\text{m})$.

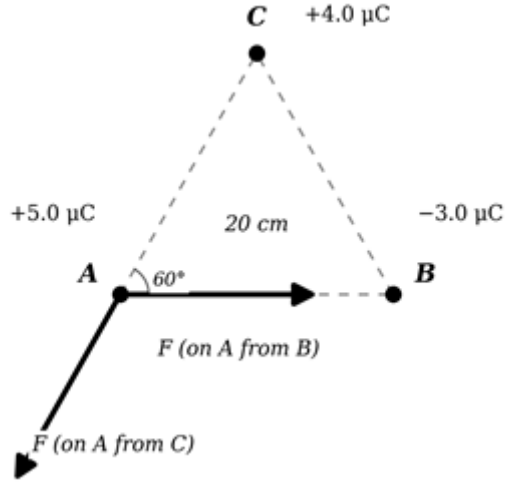


Figure: *Equilateral triangle ABC of side 20 cm with the charges marked at each vertex. The two Coulomb forces on the charge at A are shown: the force from B is attractive and points from A toward B; the force from C is repulsive and points from C through A and beyond.*

Force on A due to B (attractive, directed from A toward B):

$$\text{F}_{\text{AB}} = \text{k}\frac{|\text{Q}_\text{A}\text{Q}_\text{B}|}{\text{r}^2} = \frac{(9.0\times10^9\text{N}\cdot\text{m}^2/\text{C}^2)\times(5.0\times10^{-6}\text{C})\times(3.0\times10^{-6}\text{C})}{(0.20\text{m})^2} = 3.375\text{N}$$

Components: $\text{F}_{\text{AB,x}} = +3.375\text{N}$, $\text{F}_{\text{AB,y}} = 0\text{N}$.

Force on A due to C (repulsive, directed from C toward A and beyond):

$$\text{F}_{\text{AC}} = \text{k}\frac{|\text{Q}_\text{A}\text{Q}_\text{C}|}{\text{r}^2} = \frac{(9.0\times10^9\text{N}\cdot\text{m}^2/\text{C}^2)\times(5.0\times10^{-6}\text{C})\times(4.0\times10^{-6}\text{C})}{(0.20\text{m})^2} = 4.5\text{N}$$

The interior angle of the equilateral triangle at A is $60°$. The line CA extended past A therefore lies in the third quadrant, making $60°$ below the negative x-axis. The repulsive force F_{AC} acts along this line, so its components are:

$$\text{F}_{\text{AC,x}} = -\text{F}_{\text{AC}}\cos60° = -(4.5\text{N})\times(0.500) = -2.25\text{N}$$
$$\text{F}_{\text{AC,y}} = -\text{F}_{\text{AC}}\sin60° = -(4.5\text{N})\times(0.866) = -3.897\text{N}$$

Resultant components on A:

$$\sum\text{F}_\text{x} = 3.375\text{N} - 2.25\text{N} = 1.125\text{N}, \quad \sum\text{F}_\text{y} = 0\text{N} - 3.897\text{N} = -3.897\text{N}$$

Magnitude:

$$\text{F} = \sqrt{(\sum\text{F}_\text{x})^2 + \left(\sum\text{F}_\text{y}\right)^2} = \sqrt{(1.125\text{N})^2 + (3.897\text{N})^2} = 4.06\text{N}$$

Direction:

$$\theta = \tan^{-1}\left(\frac{|\sum\text{F}_\text{y}|}{|\sum\text{F}_\text{x}|}\right) = \tan^{-1}\left(\frac{3.897}{1.125}\right) = 73.9°$$

The resultant force on A has magnitude 4.06N, directed at $73.9°$ below the line AB, on the side of AB opposite to C.

Example 37

(a) A student claims that the electric field on the axis of a uniformly charged ring is greatest at the centre of the ring, "because the charges are closest there". Identify the error in the student's reasoning, and state the actual location on the axis at which the field is greatest.

(b) A thin straight rod of length 1.0m carries a uniform positive linear charge density $\lambda = 4.0\text{nC/m}$. Determine the electric field strength at a point P located perpendicular to the midpoint of the rod, 0.30m from it. Compare your result with the value obtained if the rod were treated as effectively infinite at the same perpendicular distance.

Solution

(a) The error is that the student weighs each charge element only by its distance, ignoring the directions of the field contributions. At the centre of the ring every element's contribution to the field at the centre is radial; by symmetry these contributions cancel in opposite pairs across the centre, and the net field at the centre is zero. The field on the axis is zero at the centre, rises to a maximum at $x = R/\sqrt{2}$ from the centre (where R is the radius of the ring) (and decays as kQ/x^2 far from the ring).

(b) Half-length $a = 0.50\text{m}$; perpendicular distance $x = 0.30\text{m}$. Then:

$$\sqrt{x^2 + a^2} = \sqrt{(0.30\text{m})^2 + (0.50\text{m})^2} = 0.583\text{m}$$

Using the finite-line field formula at the perpendicular bisector of the rod:

$$E = \frac{2k\lambda a}{x\sqrt{x^2 + a^2}}$$

$$E = \frac{2 \times (9.0 \times 10^9 \text{N} \cdot \text{m}^2/\text{C}^2) \times (4.0 \times 10^{-9}\text{C/m}) \times (0.50\text{m})}{(0.30\text{m}) \times (0.583\text{m})} \approx 206\text{N/C}$$

For the infinite-line approximation:

$$E_\infty = \frac{2k\lambda}{x} = \frac{2 \times (9.0 \times 10^9 \text{Nm}^2/\text{C}^2) \times (4.0 \times 10^{-9}\text{C/m})}{(0.30\text{m})} = 240\text{N/C}$$

The finite-line value is about $206/240 \approx 86\%$ of the infinite-line value, because the two ends of the actual rod contribute less than an extension to infinity would.

Example 38

(a) Explain why the electric potential at every point on the perpendicular bisector of an electric dipole is zero, while the electric field at the same points is not.

(b) A thin straight rod of length 1.0m carries a uniform positive linear charge density $\lambda = 4.0\text{nC/m}$. Determine the electric potential at a point P located perpendicular to the midpoint of the rod, 0.30m from it, taking the reference of zero potential at infinity.

Solution

(a) The electric potential at a point is the scalar sum of the potential contributions from each charge. The two charges of the dipole, equal in magnitude and opposite in sign, sit equidistant from any point on the perpendicular bisector; their potential contributions are equal and opposite, and cancel to zero. The electric field is a vector sum. The two contributions at a point on the bisector are vectors of equal magnitude pointing in different directions (one away from the positive charge, one toward the negative one). They do not cancel; they combine into a non-zero resultant lying parallel to the dipole axis.

(b) Half-length $a = 0.50\text{m}$; perpendicular distance $x = 0.30\text{m}$. $\sqrt{x^2 + a^2} = 0.583\text{m}$.
Using the finite-line potential formula at the perpendicular bisector:

$$V = k\lambda \ln\left(\frac{\sqrt{x^2 + a^2} + a}{\sqrt{x^2 + a^2} - a}\right)$$

$$V = (9.0 \times 10^9 \text{Nm}^2/\text{C}^2) \times (4.0 \times 10^{-9}\text{C/m}) \times \ln\left(\frac{0.583\text{m} + 0.500\text{m}}{0.583\text{m} - 0.500\text{m}}\right) = 92.5\text{V}$$

The electric potential at P is approximately 92.5V, taking the reference at infinity.

Example 39

(a) Two identical conducting spheres of mass m are suspended from a common point by insulating threads of equal length L, and each carries a charge Q. Show that at small-angle equilibrium the centre-to-centre separation of the spheres satisfies

$$d^3 = \frac{Q^2 L}{2\pi\varepsilon_0 \, mg}$$

(b) A hollow conducting sphere of inner radius 6.0cm and outer radius 10cm carries a total charge of $+12\text{nC}$ in electrostatic equilibrium. Determine the magnitude of the electric field at each of the following points: (i) the centre of the cavity, (ii) a point 3.0cm from the centre, (iii) a point 8.0cm from the centre, (iv) a point 15cm from the centre.

Solution

(a) Each sphere experiences three forces in equilibrium: weight mg downward, tension T along the thread, and Coulomb repulsion $F = kQ^2/d^2$ horizontal. Resolving the tension along the thread:

$$T\cos\theta = mg, \quad T\sin\theta = \frac{kQ^2}{d^2}$$

Dividing:

$$\tan\theta = \frac{kQ^2}{d^2 \, mg}$$

For small angles, $\tan\theta \approx \sin\theta \approx d/(2L)$. Therefore:

$$\frac{d}{2L} = \frac{kQ^2}{d^2\,mg}$$

$$\mathbf{d^3} = \frac{2L \times kQ^2}{mg} = \frac{2Q^2L}{4\pi\varepsilon_0\,mg} = \frac{\mathbf{Q^2L}}{\mathbf{2\pi\varepsilon_0\,mg}}$$

which is the required result.

(b) Let $R_1 = 0.060\text{m}$, $R_2 = 0.100\text{m}$, $Q = 12 \times 10^{-9}\text{C}$.

(i) The centre of the cavity sits at $r = 0 < R_1$; the Gaussian sphere of radius r encloses no charge, so:

$$E = 0$$

(ii) At $r = 0.030\text{m}$, still in the cavity, $r < R_1$, so $E = 0$.

(iii) At $r = 0.080\text{m}$, the Gaussian sphere lies entirely within the metal of the shell. The field inside a conductor in electrostatic equilibrium is zero, so $E = 0$.

(iv) At $r = 0.150\text{m} > R_2$, the Gaussian sphere encloses the whole of the charge Q:

$$E = \frac{kQ}{r^2} = \frac{(9.0 \times 10^9\text{Nm}^2/\text{C}^2) \times (12 \times 10^{-9}\text{C})}{(0.150\text{m})^2} = 4800\text{N/C}$$

Example 40

(a) Briefly explain how Millikan determined the elementary charge using charged oil drops.

(b) A hollow insulating sphere of inner radius 4.0cm and outer radius 10cm is uniformly charged through its shell volume with a total charge of $+8.0\text{nC}$. Determine the magnitude of the electric field at: (i) a point 3.0cm from the centre, (ii) a point 7.0cm from the centre, (iii) the outer surface of the sphere, (iv) a point 15cm from the centre.

Solution

(a) Millikan sprayed a fine mist of oil drops between two horizontal parallel plates. He selected an individual charged drop under a microscope and adjusted the potential difference between the plates until the upward electric force on its charge balanced its downward weight. From $QE = mg$, with E known from the plate voltage and spacing and m found from the drop's terminal velocity in the absence of the field, he calculated the charge Q on each drop. Repeated measurements yielded charges that were always whole-number multiples of one smallest amount, the elementary charge $e \approx 1.6 \times 10^{-19}\text{C}$.

(b) Let $R_1 = 0.040\text{m}$, $R_2 = 0.100\text{m}$, $Q = 8.0 \times 10^{-9}\text{C}$.

Then $R_1^3 = 6.4 \times 10^{-5}\text{m}^3$, $R_2^3 = 1.0 \times 10^{-3}\text{m}^3$, $R_2^3 - R_1^3 = 9.36 \times 10^{-4}\text{m}^3$.

(i) At $r = 0.030\text{m} < R_1$: the Gaussian sphere lies inside the empty cavity and encloses no charge, so:

$$E = 0$$

(ii) At $r = 0.070\text{m}$ the point lies within the charged shell, $R_1 < r < R_2$. Apply the shell-region formula:

$$E = \frac{kQ(r^3 - R_1^3)}{r^2(R_2^3 - R_1^3)}$$

With $r^3 = 3.43 \times 10^{-4}\text{m}^3$, $r^3 - R_1^3 = 2.79 \times 10^{-4}\text{m}^3$, $r^2 = 4.9 \times 10^{-3}\text{m}^2$:

$$E = \frac{(9.0 \times 10^9\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C}) \times (2.79 \times 10^{-4}\text{m}^3)}{(4.9 \times 10^{-3}\text{m}^2) \times (9.36 \times 10^{-4}\text{m}^3)} \approx 4380\text{N/C}$$

(iii) At the outer surface, $r = R_2 = 0.100\text{m}$, the Gaussian encloses the whole charge, so:

$$E = \frac{kQ}{R_2^2} = \frac{(9.0 \times 10^9\text{N} \cdot \text{m}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C})}{(0.100\text{m})^2} = 7200\text{N/C}$$

(iv) At $r = 0.150\text{m} > R_2$:

$$E = \frac{kQ}{r^2} = \frac{(9.0 \times 10^9\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C})}{(0.150\text{m})^2} = 3200\text{N/C}$$

Example 41

(a) Explain what is meant by the statement that electric charge is quantised, and give the numerical value of the smallest free charge in nature.

(b) An isolated solid conducting sphere of radius 2.0cm carries a total positive charge of $+8.0\text{nC}$. Determine: (i) the number of electrons that have been removed from the sphere, (ii) the electric potential at the centre of the sphere, (iii) the electric potential at its surface, (iv) the electric potential at a point 5.0cm from the centre, (v) the electric potential at a point 20cm from the centre.

Solution

(a) Electric charge is quantised in the sense that the charge on any object is always an integer multiple of one smallest indivisible amount, namely the charge on a single electron (or, equivalently, on a single proton, with opposite sign). This smallest free charge in nature is the elementary charge $e = 1.6 \times 10^{-19}\text{C}$.

(b) Let $R = 0.020\text{m}$ and $Q = 8.0 \times 10^{-9}\text{C}$.

(i) Number of electrons removed:

$$n = \frac{Q}{e} = \frac{8.0 \times 10^{-9}\text{C}}{1.6 \times 10^{-19}\text{C}} = 5.0 \times 10^{10}\text{electrons}$$

(ii) Inside a conductor in equilibrium the field is zero, so the potential is constant throughout the interior and equals its value at the surface. At the centre:

$$V = \frac{kQ}{R} = \frac{(9.0 \times 10^{9}\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C})}{0.020\text{m}} = 3600\text{V}$$

(iii) The surface sits at the same constant interior value, $V = 3600\text{V}$.

(iv) At $r = 0.050\text{m} > R$, the conducting sphere looks from outside like a point charge Q at its centre:

$$V = \frac{kQ}{r} = \frac{(9.0 \times 10^{9}\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C})}{0.050\text{m}} = 1440\text{V}$$

(v) At $r = 0.20\text{m}$:

$$V = \frac{kQ}{r} = \frac{(9.0 \times 10^{9}\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C})}{0.20\text{m}} = 360\text{V}$$

Example 42

(a) Explain conservation of charge as a named principle, and give an example of a process in which it must hold.

(b) Two large parallel conducting plates are held a short distance apart in air. Each plate carries the same sign of charge, with a uniform surface charge density $\sigma = 4.0\mu\text{C/m}^2$. Determine the magnitude and direction of the electric field in: (i) the region to the left of both plates, (ii) the gap between the plates, (iii) the region to the right of both plates. Take $\varepsilon_0 = 8.854 \times 10^{-12}\text{C}^2/(\text{N} \cdot \text{m}^2)$.

Solution

(a) Conservation of electric charge is a named principle of physics which states that the total electric charge of an isolated system is constant in time. This means charge may be transferred from one body to another or separated within a body, but it can be neither created nor destroyed. Example: when a glass rod is rubbed with silk, an equal magnitude of positive charge appears on the glass and negative charge on the silk; the total charge of the glass-and-silk system remains zero throughout the process.

(b) Each sheet alone produces a uniform field of magnitude $\frac{\sigma}{2\varepsilon_0}$ on each of its sides, directed away from the sheet (positive charge). With the rightward direction taken as positive, superpose the two contributions in each region.

(i) To the left of both plates, both single-sheet fields point to the left:

$$E_{\text{left}} = -\frac{\sigma}{2\varepsilon_0} - \frac{\sigma}{2\varepsilon_0} = -\frac{\sigma}{\varepsilon_0} = -\frac{(4.0 \times 10^{-6}\text{C/m}^2)}{\left(8.854 \times 10^{-12}\text{C}^2/(\text{Nm}^2)\right)} = -4.5 \times 10^{5}\text{N/C}$$

(ii) In the gap, the field of the left sheet points to the right (into the gap) and the field of the right sheet points to the left (also into the gap). The two contributions are equal in magnitude but opposite in direction:

$$E_{\text{gap}} = +\frac{\sigma}{2\varepsilon_0} - \frac{\sigma}{2\varepsilon_0} = 0$$

(iii) To the right of both plates, both single-sheet fields point to the right:

$$E_{\text{right}} = +\frac{\sigma}{2\varepsilon_0} + \frac{\sigma}{2\varepsilon_0} = +\frac{\sigma}{\varepsilon_0} = +\frac{(4.0 \times 10^{-6}\text{C/m}^2)}{\left(8.854 \times 10^{-12}\text{C}^2/(\text{Nm}^2)\right)} = +4.5 \times 10^{5}\text{N/C}$$

The field is $4.5 \times 10^{5}\text{N/C}$ directed away from the pair on each side, and zero in the gap between the plates.

Example 43

(a) Explain the three methods of charging a neutral conductor, naming each and stating in one sentence what happens to the charge in each case.

(b) A long straight wire of negligible thickness carries a uniform positive linear charge density of 5.0nC/m. Treating the wire as effectively infinite, determine: (i) the magnitude of the electric field at a perpendicular distance of 10cm from the axis, (ii) the potential difference between two points at perpendicular distances of 5.0cm and 20cm respectively from the axis, stating which point is at the higher potential.

Solution

(a) The methods are:

1. Charging by friction: rubbing two suitable insulators together transfers electrons from one to the other, leaving the two with equal and opposite net charges.

2. Charging by contact: a charged body touched briefly to a neutral conductor shares its charge with the conductor, and on separation each carries part of the original charge.

3. Charging by induction: a charged body brought near a neutral conductor without touching it polarises the conductor; momentarily earthing the far end drains the unwanted sign, and on disconnection the conductor is left with a net charge opposite in sign to the inducing body.

(b) (i) Field of the infinite line at perpendicular distance r:

$$E = \frac{\lambda}{2\pi\varepsilon_0 r} = \frac{2k\lambda}{r} = \frac{2 \times (9.0 \times 10^9 \text{Nm}^2/\text{C}^2) \times (5.0 \times 10^{-9} \text{C/m})}{(0.10\text{m})} = 900\text{N/C}$$

directed radially outward from the wire.

(ii) Potential difference between a point at distance r_a and one at distance r_b:

$$V_a - V_b = \frac{\lambda}{2\pi\varepsilon_0} \ln\left(\frac{r_b}{r_a}\right) = 2k\lambda \ln\left(\frac{r_b}{r_a}\right)$$

With $r_a = 0.050\text{m}$ and $r_b = 0.20\text{m}$:

$$V_a - V_b = 2 \times \left(9.0 \times \frac{10^9 \text{Nm}^2}{\text{C}^2}\right) \times \left(5.0 \times \frac{10^{-9}\text{C}}{\text{m}}\right) \times \ln\left(\frac{0.20\text{m}}{0.050\text{m}}\right) = 124.7\text{V}$$

The point at 5.0cm sits at the higher potential, since the potential of a positively charged line falls as one moves outward.

Example 44

(a) Explain why electric field lines never cross each other and never form closed loops in electrostatics.

(b) Four point charges are fixed at the corners of a square ABCD of side 1.0m, arranged so that A is at the top-left, B at the top-right, C at the bottom-right and D at the bottom-left. The charges are: at A, $+2.0\mu\text{C}$; at B, $-3.0\mu\text{C}$; at C, $+1.0\mu\text{C}$; at D, $+4.0\mu\text{C}$. Determine the magnitude and direction of the resultant electric field at the centre O of the square.

Solution

(a) Field lines never cross because the electric field at any point has one definite direction; if two lines crossed at a point, the field there would have two directions at once, which is impossible. Field lines never form closed loops in electrostatics because the electrostatic field is conservative; the work done by the field on a test charge taken around any closed path is zero, which forbids a continuous loop of field along which the work would accumulate.

(b) Each corner of the square lies on a diagonal that passes through the centre O. The distance from any corner to O is half the diagonal:

$$r = \frac{1.0\text{m}}{\sqrt{2}} \quad r^2 = \frac{(1.0\text{m})^2}{2} = 0.50\text{m}^2$$

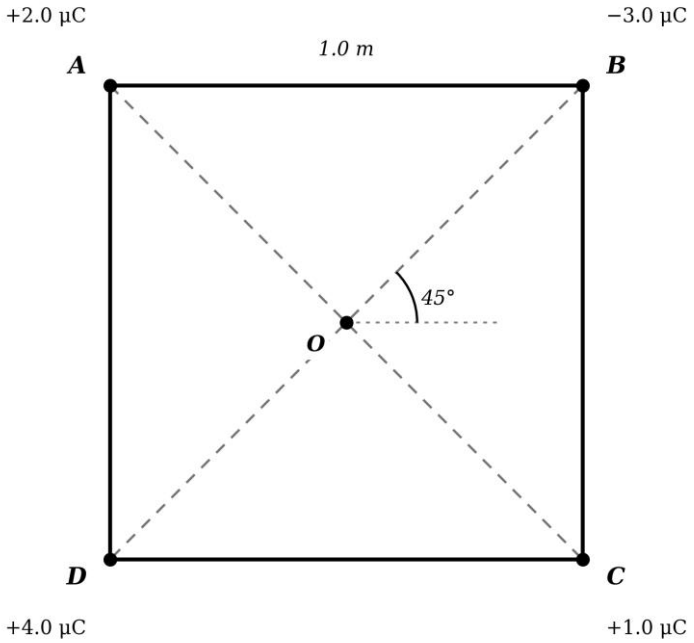


Figure: *Square ABCD with charges at each corner and centre O. Both diagonals (dashed) pass through O. Each diagonal makes 45° with the sides of the square, so every field vector at O lies along a diagonal and resolves into components by cos 45° and sin 45°.*

Field magnitudes at O due to each charge (using $E = k|Q|/r^2$):

$$E_A = \frac{(9.0 \times 10^9 \text{Nm}^2/\text{C}^2) \times (2.0 \times 10^{-6}\text{C})}{(0.50\text{m}^2)} = 3.6 \times 10^4 \text{N/C}$$

$$E_B = \frac{(9.0\times10^9\,\mathrm{Nm^2/C^2})\times(3.0\times10^{-6}\,\mathrm{C})}{(0.50\,\mathrm{m^2})} = 5.4\times10^4\,\mathrm{N/C}$$
$$E_C = \frac{(9.0\times10^9\,\mathrm{Nm^2/C^2})\times(1.0\times10^{-6}\,\mathrm{C})}{(0.50\,\mathrm{m^2})} = 1.8\times10^4\,\mathrm{N/C}$$
$$E_D = \frac{(9.0\times10^9\,\mathrm{Nm^2/C^2})\times(4.0\times10^{-6}\,\mathrm{C})}{(0.50\,\mathrm{m^2})} = 7.2\times10^4\,\mathrm{N/C}$$

Take +x along the line from A to B (parallel to side AB) and +y along the line from D to A (parallel to side AD). Each field vector at O lies along a diagonal and therefore makes $45°$ with both axes.

Field due to A. A is positive, so E_A at O points away from A, toward the lower-right corner. Its x-component is positive (rightward), its y-component negative (downward):

$$E_{A,x} = +E_A\cos45° = +(3.6\times10^4\,\mathrm{N/C})\times(0.7071) = +2.55\times10^4\,\mathrm{N/C}$$
$$E_{A,y} = -E_A\sin45° = -(3.6\times10^4\,\mathrm{N/C})\times(0.7071) = -2.55\times10^4\,\mathrm{N/C}$$

Field due to B. B is negative, so E_B at O points toward B (upper-right). Both components positive:

$$E_{B,x} = +E_B\cos45° = +(5.4\times10^4\,\mathrm{N/C})\times(0.7071) = +3.82\times10^4\,\mathrm{N/C}$$
$$E_{B,y} = +E_B\sin45° = +(5.4\times10^4\,\mathrm{N/C})\times(0.7071) = +3.82\times10^4\,\mathrm{N/C}$$

Field due to C. C is positive, so E_C at O points away from C, toward the upper-left. x-component negative, y-component positive:

$$E_{C,x} = -E_C\cos45° = -(1.8\times10^4\,\mathrm{N/C})\times(0.7071) = -1.27\times10^4\,\mathrm{N/C}$$
$$E_{C,y} = +E_C\sin45° = +(1.8\times10^4\,\mathrm{N/C})\times(0.7071) = +1.27\times10^4\,\mathrm{N/C}$$

Field due to D. D is positive, so E_D at O points away from D, toward the upper-right. Both components positive:

$$E_{D,x} = +E_D\cos45° = +(7.2\times10^4\,\mathrm{N/C})\times(0.7071) = +5.09\times10^4\,\mathrm{N/C}$$
$$E_{D,y} = +E_D\sin45° = +(7.2\times10^4\,\mathrm{N/C})\times(0.7071) = +5.09\times10^4\,\mathrm{N/C}$$

Summing components:

$$\sum E_x = (2.55 + 3.82 - 1.27 + 5.09)\times10^4\,\mathrm{N/C} = +10.18\times10^4\,\mathrm{N/C}$$
$$\sum E_y = (-2.55 + 3.82 + 1.27 + 5.09)\times10^4\,\mathrm{N/C} = +7.64\times10^4\,\mathrm{N/C}$$

Magnitude of the resultant:

$$E = \sqrt{(10.18\times10^4)^2 + (7.64\times10^4)^2}\,\mathrm{N/C} = 1.27\times10^5\,\mathrm{N/C}$$

Direction (angle above the +x axis):

$$\theta = \tan^{-1}\left(\frac{7.64}{10.18}\right) = 36.87°$$

The resultant field at O has magnitude $1.27\times10^5\,\mathrm{N/C}$, directed at $36.87°$ above the line AB, on the side of corner B.

Example 45

(a) A bird perched on a single high-voltage overhead transmission line is unharmed; a person who simultaneously touches the same line and a metal post is electrocuted. Account for this difference.

(b) Three point charges are placed at the vertices A, B and C of an equilateral triangle of side $60\,\mathrm{cm}$. The charge at A is $+5.0\,\mathrm{nC}$, at B is $-3.0\,\mathrm{nC}$, and at C is $+4.0\,\mathrm{nC}$. Calculate the electric potential at the midpoint M of side AB, taking the reference of zero potential at infinity.

Solution

(a) A bird perched on a single overhead wire rests both feet on essentially the same point of the conductor. The conductor is an equipotential, so the potential difference between the bird's feet is zero, no current flows through the bird and it is unharmed. A person standing on the ground and simultaneously touching the live wire and an earthed metal post is held at two very different potentials: the wire at tens of thousands of volts and the ground at zero. The large potential difference drives a substantial current through the person's body, with fatal consequences.

(b) M is equidistant from A and B, at a distance equal to half the side of the triangle:

$$r_{AM} = r_{BM} = \frac{0.60\,\mathrm{m}}{2} = 0.30\,\mathrm{m}$$

Vertex C lies at the apex; its distance to M is the altitude of the equilateral triangle from C:

$$r_{CM} = (0.60\,\mathrm{m})\times\frac{\sqrt{3}}{2} = (0.30\,\mathrm{m})\times\sqrt{3} = 0.520\,\mathrm{m}$$

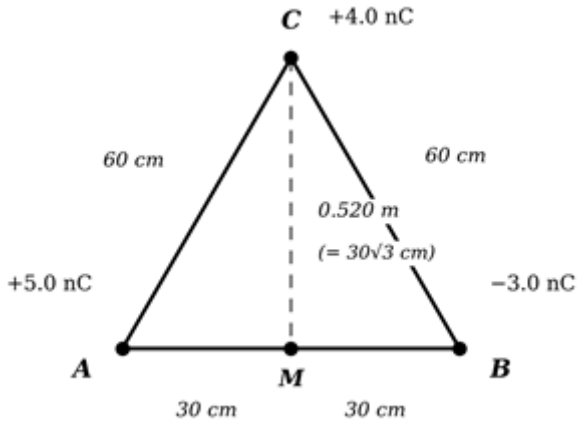


Figure: *Equilateral triangle ABC of side 60 cm. The charges at the vertices are +5.0 nC, −3.0 nC and +4.0 nC at A, B and C respectively. M is the midpoint of side AB; AM = BM = 30 cm. The altitude from C to M has length 30√3 cm = 0.520 m.*

Potential at M is the scalar sum of the three contributions:

$$V = k\left(\frac{Q_A}{r_{AM}} + \frac{Q_B}{r_{BM}} + \frac{Q_C}{r_{CM}}\right)$$

$$V = (9.0 \times 10^9 Nm^2/C^2) \times \left(\frac{(+5.0 \times 10^{-9}C)}{0.30m} + \frac{(-3.0 \times 10^{-9}C)}{0.30m} + \frac{(+4.0 \times 10^{-9}C)}{0.520m}\right) = 129.2V$$

The potential at the midpoint of AB is approximately 129V.

Example 46

(a) Cling-film stretched off the roll sometimes pulls itself onto your hand from a distance of a centimetre. Explain this observation.

(b) Two point charges $Q_1 = +5.0nC$ and $Q_2 = -5.0nC$ are fixed at two points 4.0cm apart on a straight line. A small particle of mass 8.0μg carrying a charge of $+3.0\mu C$ is released from rest at a point a on the same line, 1.0cm from Q_1. Calculate the speed of the particle when it reaches point b, located on the same line 1.0cm from Q_2. Neglect gravity.

Solution

(a) Pulling the cling-film off the roll involves friction between the film and the surface of the roll, which separates charge: the film comes away with a net charge of one sign and the roll keeps an equal and opposite charge. Brought close to the neutral hand, the charged film polarises the hand by induction, drawing opposite charge toward the surface of the hand nearer the film. The film's charge and the hand's nearer induced charge are of opposite sign, so they attract across the gap and pull the film onto the hand even before contact.

(b) The particle starts at point a with $r_{1a} = 0.010m$ and $r_{2a} = 0.030m$, and ends at point b with $r_{1b} = 0.030m$ and $r_{2b} = 0.010m$.

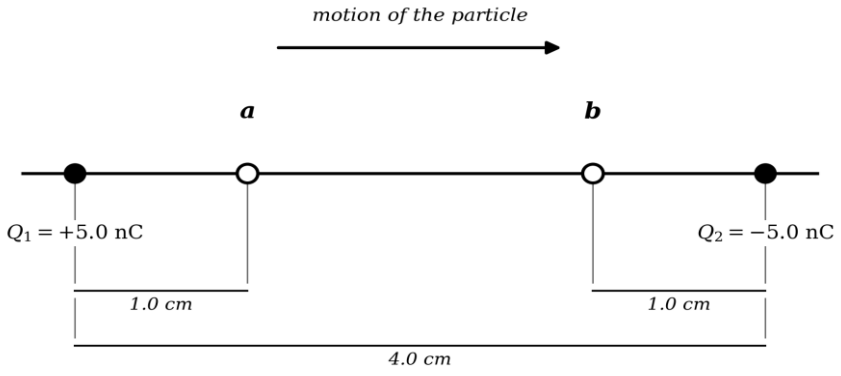


Figure: *The two fixed charges sit at the ends of a 4.0 cm line; Q_1 is at the left, Q_2 at the right. The moving particle starts at point a (1.0 cm from Q_1) and ends at point b (1.0 cm from Q_2). Open circles distinguish the trajectory endpoints from the fixed charges.*

Potential at point a:

$$V_a = k\left(\frac{Q_1}{r_{1a}} + \frac{Q_2}{r_{2a}}\right)$$

$$V_a = (9.0\times10^9\text{Nm}^2/\text{C}^2)\times\left(\frac{(+5.0\times10^{-9}\text{C})}{0.010\text{m}}+\frac{(-5.0\times10^{-9}\text{C})}{0.030\text{m}}\right) = +3000\text{V}$$

Potential at point b (by symmetry of the geometry):

$$V_b = (9.0\times10^9\text{Nm}^2/\text{C}^2)\times\left(\frac{(+5.0\times10^{-9}\text{C})}{0.030\text{m}}+\frac{(-5.0\times10^{-9}\text{C})}{0.010\text{m}}\right) = -3000\text{V}$$

Work done by the electric field on the particle from a to b:

$$W = q(V_a - V_b) = (3.0\times10^{-6}\text{C})\times\big(3000\text{V} - (-3000\text{V})\big) = 0.018\text{J}$$

By the work-energy theorem, this work equals the kinetic energy gained, since the particle starts at rest:

$$\frac{1}{2}mv^2 = W = 0.018\text{J}$$

$$v = \sqrt{\frac{2W}{m}} = \sqrt{\frac{2\times(0.018\text{J})}{(8.0\times10^{-9}\text{kg})}} = 2.1\times10^3\text{m/s}$$

The particle reaches point b with a speed of approximately $2.1\times10^3\text{m/s}$.

Example 47

(a) A petrol tanker truck is required by law to drag a metal chain along the road as it moves. Explain why the chain is essential.

(b) An electron, initially at rest, is accelerated through a potential difference of 300V in vacuum. Determine: (i) the kinetic energy acquired by the electron, (ii) its speed at the end of the acceleration. Take the electron mass $m_e = 9.11\times10^{-31}\text{kg}$ and the elementary charge $e = 1.6\times10^{-19}\text{C}$.

Solution

(a) As the tanker travels, friction between the tyres and the road, and between the moving petrol and the metal walls of the tank, separates charge and leaves the tanker with a net accumulated charge. If this charge is allowed to build up unchecked, the tanker reaches a potential at which a spark can jump between it and a nearby earthed object, such as a petrol pump or the ground, at the moment of stopping. Petrol vapour around the tanker is highly flammable, and a spark would ignite it. The trailing metal chain provides a continuous conducting path from the tanker to the ground, draining the accumulated charge to earth as fast as friction deposits it, so the tanker stays essentially at ground potential and no spark can form.

(b) (i) Work done by the field on the electron as it is accelerated through $\Delta V = 300\text{V}$:

$$W = e\Delta V = (1.6\times10^{-19}\text{C})\times(300\text{V}) = 4.8\times10^{-17}\text{J}$$

Since the electron starts at rest, all this work appears as kinetic energy:

$$\text{KE} = W = 4.8\times10^{-17}\text{J}$$

(ii) From $\text{KE} = \frac{1}{2}m_ev^2$:

$$v = \sqrt{\frac{2\text{KE}}{m_e}} = \sqrt{\frac{2\times(4.8\times10^{-17}\text{J})}{(9.11\times10^{-31}\text{kg})}} = 1.0\times10^7\text{m/s}$$

The electron acquires a kinetic energy of $4.8\times10^{-17}\text{J}$ and reaches a speed of approximately $1.0\times10^7\text{m/s}$ (about 3.3% of the speed of light).

Example 48

(a) Explain why the factor $\frac{1}{4\pi\varepsilon_0}$ appears in Coulomb's law in SI units.

(b) Two equal positive point charges, each $Q = 8.0\text{nC}$, are fixed in vacuum at the points $x = -a$ and $x = +a$ on a horizontal line, with $a = 4.0\text{cm}$ and the origin O at the midpoint between them. A small bead of mass $m = 1.0\text{g}$, carrying a positive charge $q = +8.0\text{nC}$, is threaded on a smooth insulating wire that lies along the x-axis between the two fixed charges and can slide along the wire without friction. (i) Show that the bead is in equilibrium at O. (ii) Show that for small displacements x along the wire from O, the bead executes simple harmonic motion, and derive an expression for the period of small oscillations. (iii) Calculate the period numerically.

Solution

(a) The factor $\frac{1}{4\pi\varepsilon_0}$ appears in Coulomb's law in SI units because the SI system is built upon Gauss's law as the defining relation between charge and field. By Gauss's law, the total electric flux from a point charge Q through any closed surface that encloses it is Q/ε_0; spreading this flux uniformly over a sphere of surface area $4\pi r^2$ gives $E = Q/(4\pi\varepsilon_0 r^2)$. The 4π is the geometric factor for the surface area of a sphere, and the $1/\varepsilon_0$ is the dimensional factor that converts a charge in coulombs into a field in V/m. In Gaussian (CGS) units the constant is absorbed into the definition of the unit of charge and Coulomb's law reduces to $F = \frac{q_1q_2}{r^2}$; the explicit $\frac{1}{4\pi\varepsilon_0}$ of SI is the cost of having defined the coulomb as the fundamental unit of charge.

(b) (i) Place +x along the wire, with O at the origin. The bead at position x (with $|x| < a$) feels two repulsive Coulomb forces: from the charge at $-a$, a force in the +x direction with magnitude $kQq/(x + a)^2$; from the charge at $+a$, a force in the $-x$ direction with magnitude $kQq/(a - x)^2$. The net force, with +x positive, is:

$$F(x) = kQq\left(\frac{1}{(x + a)^2} - \frac{1}{(a - x)^2}\right)$$

At $x = 0$, $F(0) = kQq\left[\frac{1}{a^2} - \frac{1}{a^2}\right] = 0$, so the midpoint is an equilibrium position.

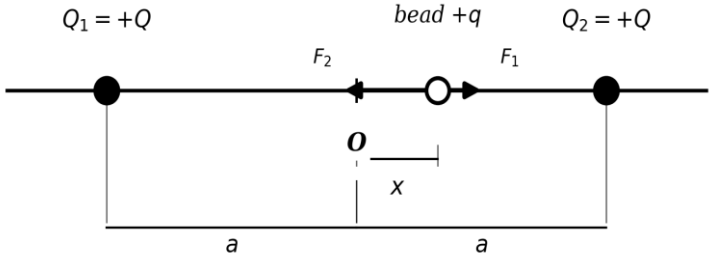


Figure: *The bead +q sits on a smooth insulating wire between the two fixed charges $Q_1 = Q_2 = +Q$. At a small displacement x to the right of O, the rightward force F_1 from the more distant charge Q_1 at $-a$ is smaller than the leftward force F_2 from the nearer charge Q_2 at $+a$, so the net force on the bead is restoring (toward O).*

(ii) For small x (with $|x| \ll a$), expand each denominator to first order in x/a:

$$\frac{1}{(x + a)^2} \approx \frac{1}{a^2}\left(1 - \frac{2x}{a}\right), \quad \frac{1}{(a - x)^2} \approx \frac{1}{a^2}\left(1 + \frac{2x}{a}\right)$$

Substituting:

$$F(x) = \frac{kQq}{a^2}\left(1 - \frac{2x}{a} - 1 - \frac{2x}{a}\right) = -\frac{4kQq}{a^3}x$$

This has the form $F = -k_{eff}\,x$, with effective spring constant $k_{eff} = \frac{4kQq}{a^3}$

By Newton's second law, $F = ma = m\frac{d^2x}{dt^2} = -k_{eff}\,x$, which is the equation of simple harmonic motion with angular frequency:

$$\omega^2 = \frac{k_{eff}}{m} = \frac{4kQq}{ma^3}$$

Period of small oscillations:

$$T = \frac{2\pi}{\omega} = 2\pi\sqrt{\frac{ma^3}{4kQq}} = \pi\sqrt{\frac{ma^3}{kQq}}$$

(iii) Substitute the numerical values:

$$T = \pi\sqrt{\frac{(1.0 \times 10^{-3}\text{kg}) \times (0.040\text{m})^3}{(9.0 \times 10^9\text{Nm}^2/\text{C}^2) \times (8.0 \times 10^{-9}\text{C}) \times (8.0 \times 10^{-9}\text{C})}} = 1.05\text{s}$$

The bead oscillates about O with a period of approximately 1.05s.

Example 49

(a) Explain why two independent natural light sources cannot produce a sustained spark-pattern in the way that two charged spheres can produce a sustained Coulomb force.

(b) An electron travelling horizontally at $v = 2.0 \times 10^7\text{m/s}$ enters the region between two horizontal parallel metal plates. The plates are $L = 4.0\text{cm}$ long and $d = 2.0\text{cm}$ apart, and a constant potential difference of 200V is maintained between them, with the upper plate positive. The electron enters midway between the plates, parallel to them, and the field outside the plate region is negligible. Take $m_e = 9.11 \times 10^{-31}\text{kg}$ and $e = 1.6 \times 10^{-19}\text{C}$. Determine: (i) the magnitude and direction of the electric field between the plates, (ii) the magnitude and direction of the electric force on the electron, (iii) the magnitude of its acceleration, (iv) the time spent between the plates, (v) the vertical deflection of the electron at the moment it emerges from the plates, and verify that it does not strike either plate during transit, (vi) the angle its velocity makes with the horizontal as it emerges.

Solution

(a) Light from two ordinary independent sources consists of wave-trains whose phases vary randomly and very rapidly with respect to one another, much faster than the eye or any practical detector can resolve. At each instant the two waves produce

some interference pattern, but the pattern shifts position with every phase fluctuation and is washed out in the time-average; no sustained pattern is observed unless the sources are phase-locked, i.e., coherent. The electrostatic interaction between two charged spheres is not a wave phenomenon at all: it is a static configuration of field vectors in space, fixed entirely by the positions and magnitudes of the charges, with no phase relationship to vary in time. The contrast is the basic difference between an electrostatic field, which is well-defined at every point in space, and an oscillating wave field, where steady patterns require coherence between sources.

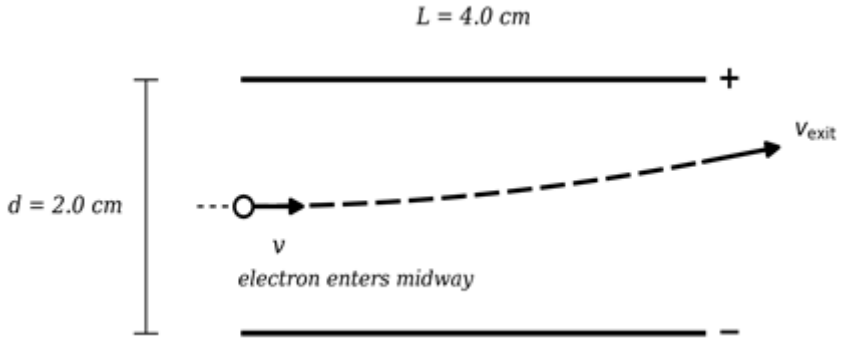


Figure: *The electron enters midway between the plates moving horizontally with speed v. The uniform electric field in the gap points downward (from + to −), so the force on the negatively charged electron is upward. The trajectory inside the plates is a parabolic curve toward the upper (positive) plate. On exit, the velocity is inclined at an angle θ above the horizontal.*

(b) (i) Field strength between the plates:

$$\mathrm{E} = \frac{\Delta \mathrm{V}}{\mathrm{d}} = \frac{(200\mathrm{V})}{(0.020\mathrm{m})} = 1.0 \times 10^{4}\mathrm{V/m}$$

directed from the upper (+) plate to the lower (−), i.e., vertically downward.

(ii) The electron carries charge $-\mathrm{e}$; the electric force is $\vec{\mathrm{F}} = (-\mathrm{e})\vec{\mathrm{E}}$, opposite to E, i.e., vertically upward (toward the positive plate). Magnitude:

$$\mathrm{F} = \mathrm{eE} = (1.6 \times 10^{-19}\mathrm{C}) \times (1.0 \times 10^{4}\mathrm{V/m}) = 1.6 \times 10^{-15}\mathrm{N}$$

(iii) Acceleration:

$$\mathrm{a} = \frac{\mathrm{F}}{\mathrm{m_e}} = \frac{(1.6 \times 10^{-15}\mathrm{N})}{(9.11 \times 10^{-31}\mathrm{kg})} = 1.76 \times 10^{15}\mathrm{m/s^2}$$

directed vertically upward.

(iv) Horizontal motion is uniform at v; time to traverse the length of the plates:

$$\mathrm{t} = \frac{\mathrm{L}}{\mathrm{v}} = \frac{(0.040\mathrm{m})}{(2.0 \times 10^{7}\mathrm{m/s})} = 2.0 \times 10^{-9}\mathrm{s}$$

(v) Vertical motion is uniform acceleration from rest in the vertical direction (entry parallel to the plates, $\mathrm{v_{y,0}} = 0$):

$$\mathrm{y} = \frac{1}{2}\mathrm{at^2} = \frac{1}{2} \times (1.76 \times 10^{15}\mathrm{m/s^2}) \times (2.0 \times 10^{-9}\mathrm{s})^2 = 3.5 \times 10^{-3}\mathrm{m}$$

That is, $\mathrm{y} = 3.5\mathrm{mm}$.

The electron entered midway between plates that are $\mathrm{d} = 2.0\mathrm{cm}$ apart, so the clearance to either plate was $\mathrm{d}/2 = 10\mathrm{mm}$. Since $\mathrm{y} < \mathrm{d}/2$, the electron emerges from the plates without striking either.

(vi) Vertical velocity at exit:

$$\mathrm{v_y} = \mathrm{at} = (1.76 \times 10^{15}\mathrm{m/s^2}) \times (2.0 \times 10^{-9}\mathrm{s}) = 3.51 \times 10^{6}\mathrm{m/s}$$

Angle with the horizontal:

$$\theta = \tan^{-1}\left(\frac{\mathrm{v_y}}{\mathrm{v}}\right) = \tan^{-1}\left(\frac{(3.51 \times 10^{6}\mathrm{m/s})}{(2.0 \times 10^{7}\mathrm{m/s})}\right) = 9.96^{\circ} \approx 10^{\circ}$$

The electron emerges from the plates deflected 3.5mm upward, with its velocity inclined at 10° above the horizontal.

Example 50

(a) Photocopiers and laser printers use electrostatic attraction of toner particles to a charged drum. Explain in two sentences how the toner forms a pattern that becomes the image on the paper.

(b) Three point charges $+\mathrm{q}$, $+\mathrm{q}$, and $-2\mathrm{q}$ are fixed at the vertices A, B, C of an equilateral triangle of side $\mathrm{a} = 30\mathrm{cm}$, with the $-2\mathrm{q}$ charge at vertex C. Take $\mathrm{q} = 5.0\mathrm{nC}$. (i) Determine the total electric potential energy of this configuration. (ii) A fourth charge $+\mathrm{q}$ is now brought from infinity to the centroid G of the triangle. State, with reasoning, the work that must be done by an external agent to achieve this.

Solution

(a) A photoconductive drum is first charged uniformly, then exposed to a light pattern representing the image: where light strikes the drum the surface becomes conducting and the charge there leaks to ground, leaving the unexposed regions still charged. Charged toner particles of opposite sign are attracted only to those still-charged regions and transfer the pattern to the drum, from where the toner is transferred to paper and fused with heat.

(b) Consider the following diagram:

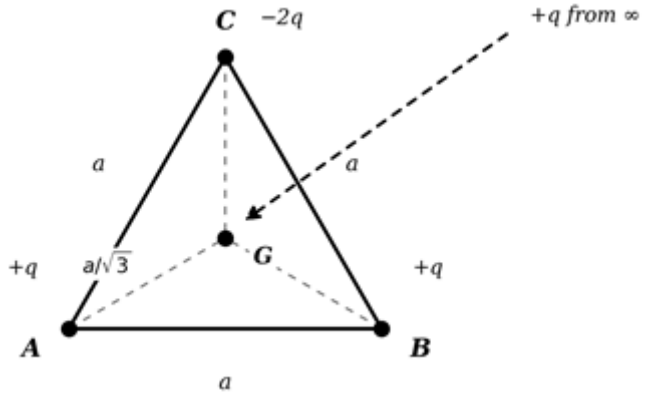


Figure: *Equilateral triangle ABC of side a. The charges at the vertices are +q, +q and −2q. The centroid G of the triangle is equidistant from all three vertices, with $r_G = a/\sqrt{3}$. The dashed arrow indicates the fourth charge being brought from infinity to G.*

(i) The total electric potential energy of a system of point charges is the scalar sum of the pair-wise potential energies $U_{ij} = kQ_iQ_j/r_{ij}$ over all distinct pairs. For three charges on an equilateral triangle of side a, all three pair separations equal a:

$$U_{AB} = \frac{k(+q)(+q)}{a} = +\frac{kq^2}{a}$$

$$U_{AC} = \frac{k(+q)(-2q)}{a} = -\frac{2kq^2}{a}$$

$$U_{BC} = \frac{k(+q)(-2q)}{a} = -\frac{2kq^2}{a}$$

Sum:

$$U = U_{AB} + U_{AC} + U_{BC} = \frac{kq^2}{a} - \frac{2kq^2}{a} - \frac{2kq^2}{a} = -\frac{3kq^2}{a}$$

Numerically:

$$U = -3 \times \frac{(9.0 \times 10^9 \text{Nm}^2/\text{C}^2) \times (5.0 \times 10^{-9}\text{C})^2}{(0.30\text{m})} = -2.25 \times 10^{-6}\text{J} = -2.25\mu\text{J}$$

The negative value tells us the system is bound: an external agent would have to do positive work of $2.25\mu\text{J}$ against the net electrostatic forces to disassemble the configuration to infinite mutual separation.

(ii) The work an external agent must do in bringing a charge q' from infinity to a point P is $W = q'V_P$, where V_P is the electric potential at P due to all other charges already in place. The centroid G of an equilateral triangle is equidistant from all three vertices, at distance:

$$r_G = \frac{a}{\sqrt{3}} = \frac{0.30\text{m}}{\sqrt{3}} = 0.173\text{m}$$

Potential at G due to the three vertex charges:

$$V_G = \frac{k}{r_G}\big((+q) + (+q) + (-2q)\big) = \frac{k}{r_G} \times 0 = 0$$

The contributions from the two $+q$ charges and the one $-2q$ charge cancel exactly at the centroid because the three distances are equal and the charges sum to zero. Therefore:

$$W = qV_G = q \times 0 = 0\text{J}$$

No work is done by the external agent in bringing the fourth charge from infinity to the centroid.

The result is counter-intuitive (three nearby charges and yet zero work), but is a clean consequence of the equidistance of G from the three vertices and the fact that the three charges sum to zero.

From the basic fact that charge exists only in discrete quanta, through Coulomb's law and the field it generates, through Gauss's law applied to lines, sheets, and spheres, to the electric potential and the work-energy bridge $W = q\Delta V$, the chapter has built electrostatics step by step, and the fifteen miscellaneous examples have exercised those ideas in combination. With the topic now viewed as a whole, we are well prepared to move on to the Digging Deeper Exercise on the next page.

DIGGING DEEPER EXERCISE 10

EXERCISE 10A: BINDER QUESTIONS

Question 1

When two or more point charges are present, the resultant electric field at a point in space is found by vector addition of the contributions of the individual charges, but the resultant electric potential at the same point is found by scalar (arithmetic) addition. Explain why the two quantities are combined by different rules.

Question 2

Explain in your own words what an electric field is, and how its strength at a point is operationally measured.

Question 3

On the surface of an irregularly shaped charged conductor in electrostatic equilibrium, the surface charge density is greater at regions of high curvature, such as sharp points and edges, than at regions of low curvature, such as broad flat surfaces. Account for this concentration of charge.

Question 4

Explain why the electric field inside an empty cavity of a hollow charged conductor in electrostatic equilibrium is zero, even though the surface of the conductor carries a non-zero charge.

Question 5

A person who has walked across a dry carpet sometimes draws a small visible spark when reaching for a metal doorknob. Explain how this is consistent with the idea that air has a finite dielectric strength.

Question 6

Provide the statement of Gauss's law for electric fields, and explain in one or two sentences why the law is particularly useful for finding the field of a charge distribution that possesses a high degree of symmetry.

Question 7

Explain what is meant by an equipotential surface, and explain three of its properties.

Question 8

Explain why no work is done by the electric field in moving a test charge along an equipotential surface, by giving two independent arguments (one in terms of the potential difference, and one in terms of the directions of the field and of the displacement).

Question 9

The relationship between the electric field strength and the electric potential along a chosen direction is $\mathrm{E} = -\frac{\mathrm{dV}}{\mathrm{dr}}$. Explain this relationship in words, and explain what the negative sign means physically.

Question 10

A conductor connected to the Earth by a wire is said to be earthed, and is assigned a potential of zero. Explain (i) why the Earth is treated as the universal reference of zero potential, and (ii) what happens physically to the charge on a previously charged conductor at the moment it is earthed.

EXERCISE 10B: REAL QUESTIONS

Question 11

Kipanga: *"I noticed that the lightning during last night's storm struck the tall coconut tree behind our house in Mtwara, not the shorter mango tree right next to it."*

Mr. Akilikubwa: *"You have noticed something every lightning conductor in the world is designed around."*

Explain why a lightning bolt is far more likely to strike a tall pointed object than a shorter one standing close beside it.

Question 12

Kipute pulls a wool jumper over her head on a dry day and her hair stands out as straight wisps that refuse to lie flat.

Kipanga: *"The hairs are pushing each other apart."*

Mr. Akilikubwa: *"You are half right."*

Explain Kipanga's correct point, and the second mechanism that Mr. Akilikubwa is adding.

Question 13

A lightning conductor on a tall building has two unmistakable features: a sharp metal point projecting above the roof, and a thick copper strap leading from the point down to a heavy metal plate buried deep in moist soil at the building's foot. Explain the purpose of each of the two features.

Question 14

A car parked in the open during a thunderstorm is sometimes struck by lightning, yet the people sitting inside it are unharmed. Explain why the metal body of the car protects them.

Question 15

An inkjet printer forms each letter on the page by firing a stream of tiny ink droplets between a pair of charged metal plates. Each droplet is given a controlled amount of charge as it leaves the nozzle, and the print head varies that charge from droplet to droplet. Explain how varying the charge on each droplet determines where on the page it lands.

Question 16

At the Tanga Cement factory, the hot flue gas leaving the kiln passes through an electrostatic precipitator before reaching the chimney. Inside the precipitator, a very high voltage is applied between a central wire and the surrounding earthed collecting plates, and the dust and ash particles are removed from the gas. Explain how the precipitator removes the dust.

Question 17

At an airport, the operating manual for a fuel truck instructs the operator to clip a heavy bonding cable from the truck to the body of the aircraft before opening the fuel valve, and to remove the cable only after the valve has been closed. Explain why the bonding cable is essential, even when the operator is standing on dry tarmac in rubber boots.

Question 18

When lightning strikes the ground close to the foot of a tall isolated tree, a person sheltering beneath that tree may be killed by the strike even though the lightning did not pass through the body. Explain why the ground itself becomes dangerous near the point at which the lightning enters it.

Question 19

A plastic comb rubbed briskly on the hair retains its charge for many minutes and continues to attract small pieces of paper for a long time. A clean metal rod, held in the bare hand and rubbed in exactly the same way on the same hair, cannot be made to attract paper at all. Explain the difference.

Question 20

In a school physics demonstration in Mwanza, the teacher holds one end of a long fluorescent tube near the highly charged metal dome of a Van de Graaff generator without making contact, and without connecting the tube to any electrical supply. The tube glows along its length with a faint flickering light. Explain how the tube can glow without being plugged in.

EXERCISE 10C: HOT QUESTIONS

Question 21

A point charge of $+4.0\mu C$ is placed at the origin of an x-axis, and a second point charge of $-2.0\mu C$ is placed at $x = 0.30m$. Determine the position on the x-axis at which the net electric field due to the two charges is zero.

Question 22

Four point charges of magnitudes $+2.0\mu C$, $-3.0\mu C$, $+2.0\mu C$ and $-3.0\mu C$ are placed in that cyclic order at the four corners of a square of side 0.10m. Find the magnitude of the net electric field at the centre of the square.

Question 23

In the right-angled triangle ABC, the right angle is at vertex B, with AB = 3.0cm and BC = 4.0cm. A point charge of +6.0nC is fixed at A and a point charge of +8.0nC is fixed at C. Find the magnitude and direction of the resultant electric field at vertex B.

Question 24

Identical point charges of magnitude +Q are placed along the x-axis at the positions x = 1.0m, 2.0m, 4.0m, 8.0m, and so on, with each successive charge at twice the distance of its predecessor from the origin. Find the magnitude of the electric field at the origin in terms of kQ.

Question 25

A uniformly charged conducting sphere of diameter 0.60m carries a uniform surface charge density of $\sigma = 50\mu C/m^2$. Determine (i) the magnitude of the electric field just outside the surface of the sphere, (ii) the total charge on the sphere, and (iii) the total electric flux through any closed surface that completely encloses the sphere.

Question 26

A student calculates the electric field at a point inside a uniformly charged insulating solid sphere of total charge Q and radius R, at a perpendicular distance r from the centre (with r < R), by writing $E = kQ/r^2$. Identify the conceptual error in the student's reasoning, and derive the correct expression for the magnitude of the electric field inside the sphere at radius r.

Question 27

Two large parallel sheets carry uniform surface charge densities of $+\sigma$ and $-\sigma$ respectively, with $\sigma = 6.0\mu C/m^2$. Find the magnitude of the electric field (i) in the region between the two sheets, and (ii) in the region outside the sheets.

Question 28

Three point charges of magnitudes Q, +q and +q are placed at the vertices of an equilateral triangle of side L. Given that the total electrostatic potential energy of the three-charge system is zero, find Q in terms of q.

Question 29

Two thin concentric spherical shells of radii r and R (with r < R) each carry the same uniform surface charge density σ. Find the electric potential at the common centre of the two shells, taking the potential to be zero at infinity.

Question 30

Two equal positive charges, each of magnitude Q, are fixed on the x-axis at the positions x = +a and x = −a. A small particle of charge +q and mass m is placed at the midpoint between them, x = 0. The particle is given a small displacement x along the x-axis and released. (a) Show that the resulting motion is simple harmonic. (b) Find the angular frequency ω in terms of Q, q, m, a, and Coulomb's constant k.

Question 31

A student computes the work done by the electrostatic field in moving a charge of 3.0C along a closed loop that lies entirely on an equipotential surface of value 12V. The student writes $W = qV = (3.0C)(12V) = 36J$. Give the correct value of the work done, and identify the conceptual error in the student's reasoning.

Question 32

Two horizontal parallel plates are placed 8.0mm apart and are maintained at a potential difference of 980V, with the upper plate at the higher potential. A small negatively charged oil drop of mass 1.0×10^{-14}kg is observed to be at rest between the plates. Calculate the number of excess electrons on the drop. Take $g = 9.8m/s^2$ and $e = 1.6 \times 10^{-19}C$.

Question 33

Two identical small pith balls, each of mass m and carrying the same charge Q, are suspended by silk threads of equal length L from a common point. In air the two threads come to rest making an angle of 30° between them. When the entire system is then immersed in a non-conducting oil of relative permittivity K, the threads come to rest making an angle of 18° between them. Assuming small-angle equilibrium throughout, and that the buoyancy on the balls in the oil is negligible, find K.

Question 34

An electrostatic precipitator at a Tanzanian cement factory must remove ash particles of diameter 1.0μm from the flue gas. Each particle is assumed to acquire a charge of 100e by collision with ions close to the central wire, and must then be driven sideways across a 15cm gap onto an earthed collecting plate during the 2.0s residence time of the gas inside the precipitator. Take the viscosity of the flue gas as that of air, $\eta = 1.8 \times 10^{-5}$Pa·s. Estimate the minimum applied electric field that will deliver every particle onto the plate within the residence time, and judge whether the chosen field is realistic in light of the dielectric strength of dry air, 3.0×10^6V/m.

Question 35

A laboratory Van de Graaff generator has a spherical metal dome of diameter 12cm and is rated to operate at a potential of 350kV above ground. Calculate the electric field at the surface of the dome at the rated potential, and determine whether the dome can hold this potential in ordinary dry air without breakdown. If the rated potential cannot be held, suggest two design modifications that would allow the dome to operate at 350kV safely. Take the dielectric strength of dry air as 3.0×10^6V/m.

Chapter 11

CAPACITANCE

INTRODUCTION

The old black-and-white television in the Miono physics prep room had been unplugged for two days. The lab technician, his hands full of soldering iron and replacement parts, glanced up from his workbench at the corner of the room. He was waiting for a new picture tube. The cabinet of the old set sat with its back panel removed, exposed to the morning light.

Kipanga, never one to leave equipment unexplored, stepped over to look inside. His curiosity rarely consulted his sense of caution. He reached his hand in past the chassis, intent on examining the picture tube. The technician, glancing up just a moment too late, opened his mouth to call out.

SNAP!

Kipanga yanked his hand back, sucking at the side of his thumb. A small red mark had risen on it. He stared at the back of the television, then at his hand, then at the technician, with all the affront of someone who has been personally insulted by an inanimate object.

Kipanga: *It bit me! The television bit me, and it is not even plugged in!*

Mr. Akilikubwa, who had been writing up notes at the doorway of the prep room, looked up with the slow, settled expression of a man who has seen this exact scene before, many times, and resigned himself to its repetition.

Mr. Akilikubwa (without rising from his notebook)**:** *Disconnected is not the same as discharged, Kipanga. Inside an old television there is a part called a capacitor, and in fact there are many such parts, including one very large one connected to the picture tube. It is charged to thousands of volts when the set is in use. When the set is unplugged, that capacitor holds onto its charge for hours, even days, sometimes weeks. The technician knows this. That is why he never reaches inside before discharging it himself.*

Kipute (from across the lab, laughing)**:** *Sir, this is the same boy who, three weeks ago, was told not to lick the terminals of a battery to test whether they were still good.*

Kipanga: *I learn by experience. The book only tells you the answer. The shock tells you the reason.*

Mr. Akilikubwa rose, walked over to the bench, and picked up a long screwdriver with a thick insulated rubber handle. He touched the metal tip of the screwdriver to the anode cap on top of the picture tube and then immediately to the chassis of the set. A small, brief snap followed; it was smaller than the one that had stung Kipanga's hand. He held the tip in place for a moment longer, withdrew it, and repeated the action a second time. The second snap was barely audible.

Mr. Akilikubwa: *Now it is safe to work on. The first shorting drains the bulk of the stored charge. The second is for anything that might have re-built itself from the residual fields. A technician who skips this step is the kind of technician one only meets once.*

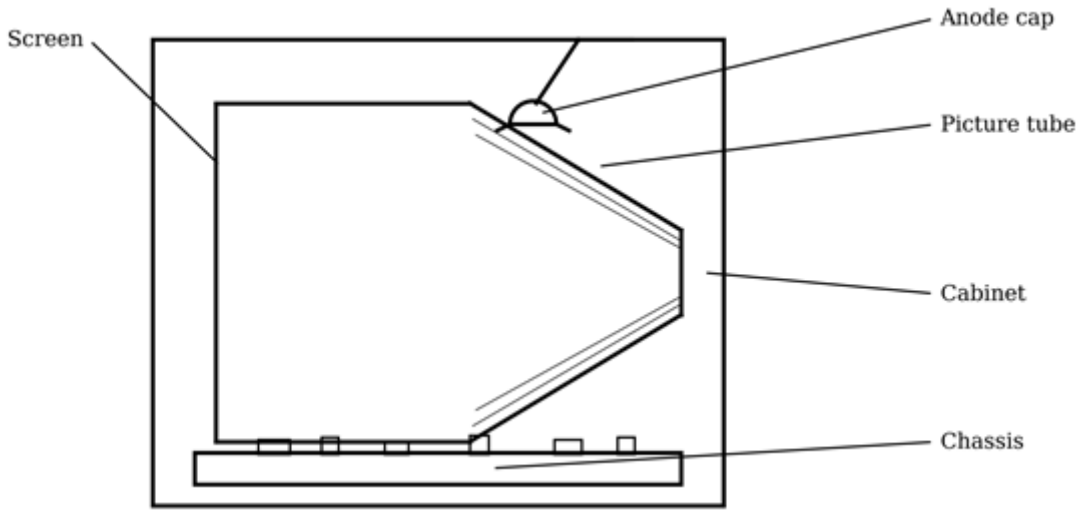


Figure: ***Cut-away view of an old CRT television set.*** *The picture tube, the chassis, and the anode cap (which holds the high voltage) are labelled. A capacitor inside the set keeps the anode cap charged for a long time after the set is unplugged.*

The capacitor is one of the three foundation components of any circuit that does something more interesting than light a single bulb. The other two are the resistor and the inductor; the resistor has been with us since O-level, and the inductor is the subject of Chapter 13. Capacitors are the components that store electric charge for later use. They make the sudden flash of a camera possible; they smooth out the trembling voltage from a power supply; they tune a radio to one station among many; they fire

the high-voltage spark plug of a petrol engine; and, as Kipanga has just discovered, they hold their charge long after their source has been removed.

In Chapter 10 we calculated the electric field between two large parallel plates carrying opposite surface charges, and found it to have the simple value $E = \frac{\sigma}{\varepsilon_0}$. In this chapter we use that result to derive a formula for how much charge a capacitor of given geometry holds for a given applied voltage. Then we extend the result to combinations of capacitors, to capacitors with insulating filler material between their plates, to the energy they store, to the force their charged plates exert on one another, and to the curve of charge versus time as a capacitor fills and empties through a resistor. We close with a set of applications that use capacitors in serious work: the camera flash, the defibrillator that revives a stopped heart, the soil-moisture sensor on a farm in Mbeya, the power-factor correction bank at a substation in Tanga, and several more.

One thing should already be clear from Kipanga's small adventure with the television: a capacitor remembers. It remembers the charge put on it long after the source has been removed. To work safely with capacitors, and to use them intelligently, we need to understand exactly how they remember, what determines how much they remember, and how they release what they remember when finally given the chance. The next eight sections of this chapter answer those questions, in the order they are best learned. Disconnected, as Mr. Akilikubwa reminded us, is not the same as discharged.

THE CAPACITOR

Open a portable radio, the controller of a desk fan, or the power-supply box of any old computer. Among the resistors, the tiny black integrated-circuit chips, and the loops of copper wire, you will see something tubular, sometimes blue, sometimes black, sometimes silver, with two wire legs sticking out of one end and the markings "10µF 25V" or "470µF 16V" printed on its side. That is an electrolytic capacitor. Pry open an old desk radio and you will find a small beige ceramic disc capacitor with a code such as "104" or "103" stamped on it. Look inside a digital voltmeter and you will see a flat rectangular box capacitor of the polyester or film type. Inside a camera, the same shape but larger, charged to several hundred volts, waits patiently to fire the flash.

These are all capacitors. They look different on the outside because of the materials and the geometries inside, but the underlying physics is the same in every case. **A capacitor** *is a device for storing electric charge on a pair of conductors held very close together but never allowed to touch.* The conductors are usually flat plates. The thin gap between them is filled with an *insulating material* called a **dielectric**, which we will discuss in detail later in the chapter. For now we will simply take the gap to be air or vacuum.

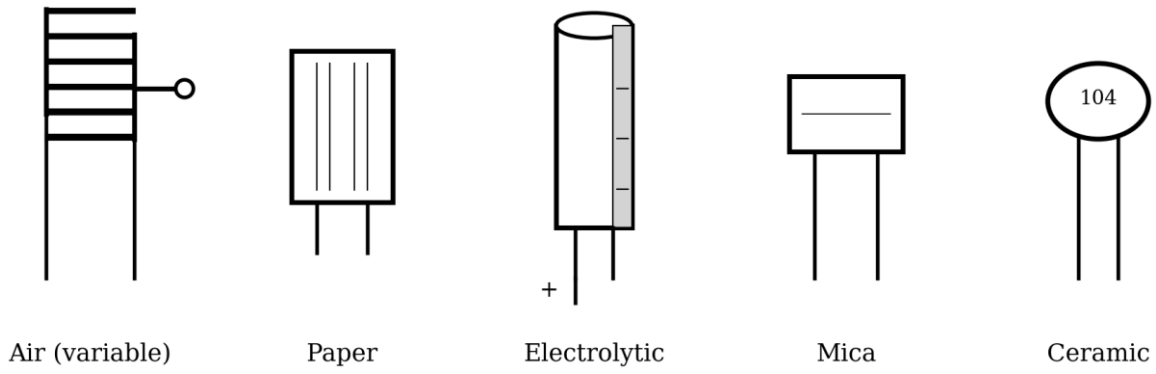


Figure: ***Five common capacitor types.*** *Each is named after the material that fills the gap between its plates.*

To get a feel for what 'storing charge' looks like in practice, look at the simple charging circuit in the figure. A battery of voltage V is connected through a switch S to a capacitor with plates labelled M and N. While the switch is open, both plates are electrically neutral. The moment the switch is closed, the battery starts pushing electrons from plate M through the external wire, around through the battery, and onto plate N, where they begin to accumulate. Plate N becomes more and more negative; plate M, which has just lost the same number of electrons, becomes equally positive.

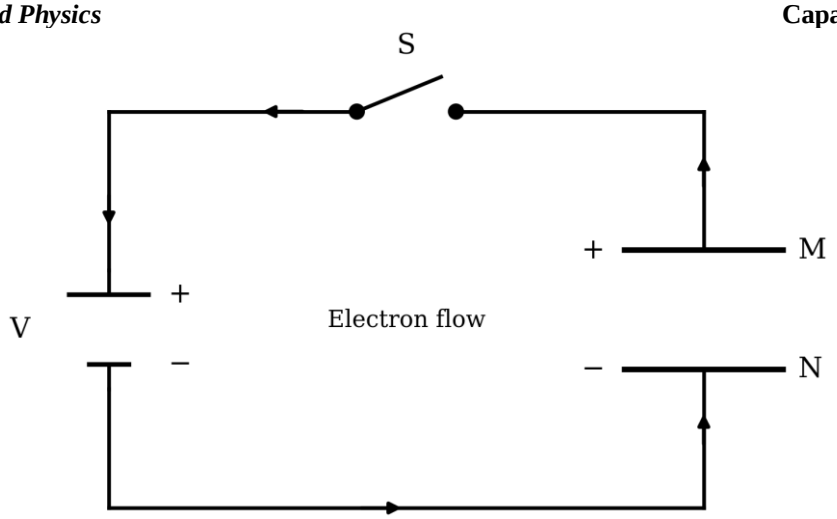


Figure: ***Charging a capacitor.*** *Electrons leave plate M, flow through the external circuit, and accumulate on plate N. The electrons never cross the gap between the plates; the dielectric prevents that.*

The electrons do not at any stage cross the gap between the plates. They cannot because the dielectric is an insulator. They only go around the outside, through the battery, depositing themselves on plate N from outside while peeling off plate M from outside. By the time the capacitor is fully charged, plate N carries a charge of magnitude Q with negative sign, and plate M carries the same magnitude Q with positive sign. The total charge on the device is $+\mathrm{Q} + (-\mathrm{Q}) = 0$. The capacitor as a whole is electrically neutral, even when it is "holding a charge of Q". When we say the capacitor holds charge Q, we mean either plate alone carries that magnitude; the total is always, at every instant, zero. The plates have simply taken sides.

How much charge piles up depends on two things only: how hard the battery pushes (its voltage V) **and** how generously the capacitor is built to receive the charge. The first is set by the supply; the second is set by the capacitor. Experiment shows that for a given capacitor, the charge Q piled on either plate is strictly proportional to the applied voltage V: double V and Q doubles, halve V and Q halves. The constant of proportionality is called the **capacitance** of the capacitor, written **C**:

$$\mathbf{C} = \frac{\mathbf{Q}}{\mathbf{V}}$$

This is the defining equation of capacitance. Read it carefully: it says that **capacitance** *is the amount of charge stored per unit of voltage applied.* A capacitor of large capacitance stores a lot of charge for each volt applied to it; a capacitor of small capacitance stores very little.

The SI unit of capacitance is the **farad (F),** in honour of Michael Faraday, the English physicist who first studied capacitors and dielectrics in serious detail in the 1830s. By definition, a capacitor of 1F stores 1C of charge when 1V is applied across its plates:

$$1\mathrm{F} = \frac{1\mathrm{C}}{1\mathrm{V}}$$

By any reasonable engineering standard, the farad is an absurdly large unit. A typical desk-radio capacitor lives in the picofarad range; six trillionths (10^{-12}) of a farad. The tubular capacitors in your power supply live in the microfarad range; millionths (10^{-6}) of a farad. Only one class of specialised device, the so-called supercapacitor or ultracapacitor, reaches whole farads, and it pays for the privilege with a low working voltage and a noticeably bulky physical size. The multiples you will routinely meet in this chapter, and in every examination paper from here onwards, are:

$$1\mathrm{mF} = 10^{-3}\mathrm{F}$$
$$1\mu\mathrm{F} = 10^{-6}\mathrm{F}$$
$$1\mathrm{nF} = 10^{-9}\mathrm{F}$$
$$1\mathrm{pF} = 10^{-12}\mathrm{F}$$

The microfarad and the picofarad do almost all the practical work. A 1F capacitor on a circuit board would be a curiosity, not a component.

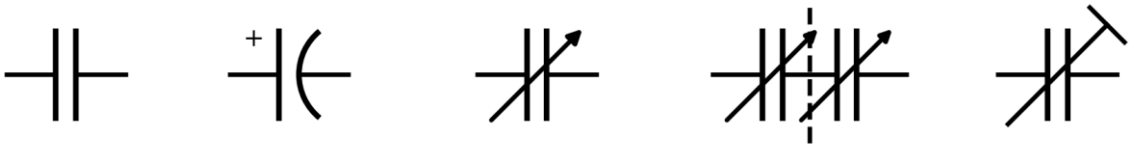
A brief word now on the five common capacitor types and where each is at home in a circuit.

1) **Air capacitors** use *ordinary air* as their dielectric. They are still common in the variable-tuning sections of older AM radios; turning the tuning knob slides one set of plates partway out from between another set, changing the effective overlap area and therefore the capacitance.
2) **Paper capacitors** use *waxed* or *oil-impregnated paper* as the dielectric. They were once the standard in radio receivers, motor-starting circuits, and old electrical machinery; modern equipment has largely replaced them with film capacitors, but they remain common in equipment that comes back to the bench for repair.
3) **Electrolytic capacitors** achieve very large capacitance values in small packages by using a *very thin chemical layer* (the *electrolyte*) as the dielectric between two metallic foils. They appear wherever a circuit needs to store a large amount of

charge: power-supply smoothing, motor-control circuits, audio amplifier coupling. They are polarised. Each electrolytic capacitor has a definite positive lead (often the longer leg, sometimes marked with stripes on the body) and a definite negative lead. Connecting an electrolytic capacitor the wrong way around results in a quiet hiss, a sharp burning smell, and a permanent change in the capacitor's circuit behaviour, in the sense that it stops being a capacitor at all.

4) **Mica capacitors** use thin *sheets of natural mica* as the dielectric between metal plates. Mica has a remarkably stable dielectric constant and a very high dielectric strength, so mica capacitors are prized for accuracy at high frequencies, particularly in radio oscillator circuits and precision measurement instruments.
5) **Ceramic capacitors** use a *thin disc of ceramic*, usually a *titanate* or *oxide compound*, as the dielectric. They are by far the cheapest and most numerous capacitors on a modern circuit board, performing routine bypass and decoupling duties wherever a little capacitance is needed. The **"104"** or **"103"** stamped on a ceramic capacitor reads as *10 followed by 4 zeros (or 3 zeros) in picofarads, giving* $100nF$ *or* $10nF$ *respectively.*

In circuit diagrams, capacitors are represented by a small standard set of symbols, shown in the following figure.



Fixed non-polarised | Fixed polarised | Single variable | Dual variable | Preset

Figure: *Standard circuit symbols for fixed (non-polarised, polarised) and variable (single, dual, preset) capacitors.*

One small geometric matter is worth dispatching before we move on. The two plates of a capacitor **need not** be the same size. If we built a capacitor with one plate a small coin and the other plate a large sheet of metal, the small coin would simply carry a higher surface charge density (charge per unit area) than the large sheet, but the total charges on the two plates would remain exactly equal in magnitude and opposite in sign. This is not a coincidence of geometry; it is the consequence of charge conservation. Every electron that arrived on the negative plate had to leave the positive plate. The battery cannot manufacture electrons out of thin air; it can only push electrons around a closed loop. Whatever number of electrons arrives at one plate must have departed the other.

What follow are three short worked examples in which nothing exciting happens. They are here because every formula in physics needs an easy outing before it is asked to do real work.

BINDER Example 1

A capacitor of capacitance $5\mu F$ carries a charge of $60\mu C$ on either plate. Find the potential difference between the plates.

Solution

From the defining equation of capacitance:

$$C = \frac{Q}{V} \qquad \Rightarrow \qquad V = \frac{Q}{C}$$

Substituting the given values:

$$V = \frac{60 \times 10^{-6} C}{5 \times 10^{-6} F} = 12V$$

The potential difference between the plates is $12V$.

Making Sense of the Answer: *Twelve volts is a perfectly ordinary working voltage for a microfarad-class capacitor; nothing here is unusual or extreme. The microfarad and the microcoulomb are companion units in the same way that the volt and the ampere are companions in everyday electrical work; both prefixes mean "one millionth", and when they appear together the result is a clean voltage of order ten volts. The 5μF–60μC–12V triple sits comfortably in the middle of the working range of any small microfarad capacitor.*

Think Like a Physicist: *Watch the prefixes cancel. Sixty divided by five is twelve, and the "micro" on top divides the "micro" on the bottom, leaving a pure voltage in volts. Whenever the charge and the capacitance are quoted in matching prefixes (microcoulombs with microfarads, nanocoulombs with nanofarads, picocoulombs with picofarads), the prefixes cancel and the answer is in straight volts. Noticing this saves several mental steps.*

BINDER Example 2

A capacitor of capacitance $90pF$ is connected across a $12V$ battery. Find the number of electrons transferred from one plate to the other during charging.

Solution

The charge that accumulates on either plate when the capacitor is fully charged is:

$$Q = CV = (90 \times 10^{-12}\text{F})(12\text{V}) = 1.08 \times 10^{-9}\text{C}$$

Each electron carries a charge of magnitude $e = 1.6 \times 10^{-19}\text{C}$. The number of electrons N transferred from one plate to the other is therefore the total charge divided by the charge per electron:

$$N = \frac{Q}{e} = \frac{1.08 \times 10^{-9}\text{C}}{1.6 \times 10^{-19}\text{C}} = 6.75 \times 10^{9}$$

The number is 6.75×10^{9} electrons (about 6.75 billion electrons move from one plate, through the external circuit, and onto the other plate during the charging process).

Making Sense of the Answer: *Six and three-quarter billion electrons sounds like an unimaginable number, but in fact it is a comparatively small number of electrons by physics standards. One coulomb of charge contains about 6.25×10^{18} electrons; what we transferred here is only about one part in a billion of a coulomb. The reason capacitors store such "small" numbers of electrons in absolute terms is that even a tiny imbalance of charge between two close plates produces a substantial voltage. Twelve volts across 90 picofarads is the comfortable operating range of any pocket radio; the same twelve volts across a much larger pair of plates might store a microcoulomb, ten thousand times what we just calculated, yet still only a millionth of a coulomb. Electric forces are strong, and small charge differences make large voltages when the plates are close.*

Think Like a Physicist: *Any "how many electrons" question hides a two-step calculation: first find the total charge using the capacitance formula, then divide by the magnitude of the electronic charge. The two steps are independent of each other and are best written on two separate lines, since trying to crush them into one over-stuffed equation invites algebra errors. Notice also the scale of the answer: a few billion electrons, not a few million or a few trillion. Once a reader has done two or three of these calculations, the order of magnitude: billions, for a small picofarad-class capacitor at battery voltages becomes a useful sanity check on every future answer.*

REAL Example 3

In the school prep room, Kipute opens a small drawer of spare components and pulls out a ceramic disc capacitor with the marking 104 stamped on its face. She wants to know how much charge the capacitor will hold when she connects it across a 5V supply. Find (a) the charge stored on either plate when the 5V supply is connected; (b) the number of electrons transferred from one plate to the other during charging.

Solution

(a) The marking 104 is read as 10 followed by 4 zeros, with the result in picofarads. The first two digits give the leading value, which is 10, and the third digit gives the number of trailing zeros, which is 4. The value of the capacitance is therefore:

$$C = 10 \times 10^{4}\text{pF} = 10^{5}\text{pF} = 10^{5} \times 10^{-12}\text{F} = 10^{-7}\text{F}$$

From the defining equation of capacitance, the charge stored on either plate when 5V is applied is:

$$Q = CV = (10^{-7}\text{F})(5\text{V}) = 5 \times 10^{-7}\text{C} = 0.5\mu\text{C}$$

(b) Each electron carries a charge of magnitude $e = 1.6 \times 10^{-19}\text{C}$. The number of electrons transferred from one plate to the other during charging is therefore:

$$N = \frac{Q}{e} = \frac{5 \times 10^{-7}\text{C}}{1.6 \times 10^{-19}\text{C}} = 3.13 \times 10^{12}\ \text{electrons}$$

Making Sense of the Answer: *Half a microcoulomb is more than the 1.08 nanocoulombs of the previous example by a factor of about five hundred. The two examples differ in capacitance by a factor of about a thousand (100nF versus 90pF) and in voltage by about a half (5V versus 12V), so the ratio of charges should be roughly 1000 ÷ 2 = 500, which is what we got. The arithmetic checks against the physics, and the habit of estimating ratios like this catches half the algebra mistakes a student can make on a capacitor question.*

Think Like a Physicist: *A three-digit code stamped on a ceramic capacitor is read in the same way the colour bands on an old resistor are read: the first two digits give the leading value, and the third digit gives the power-of-ten multiplier. The result is always in picofarads, the smallest of the common engineering units for capacitance. Knowing the convention turns a component on the bench into a usable value in a few seconds, with no datasheet required.*

That is what the formula $C = \frac{Q}{V}$ looks like with its training wheels still on. From here on it works without them.

Across the corridor from the prep room, while the technician was still waiting for the replacement television tube, Mr. Akilikubwa had set up a small display on the bench: a 1000-microfarad electrolytic capacitor, a 9V battery, a low-current LED, and a 220Ω resistor.

Mr. Akilikubwa: *Kipanga, since you have already met a charged capacitor today, you may as well be the person who charges this one. Touch the leads of this capacitor briefly across the two terminals of the battery, just a second is enough, and then place those same leads across the LED with this resistor in series. Mind which way around the LED goes. It is polarised. The longer leg is the positive lead.*

Kipanga touches the capacitor leads to the battery for about a second, then transfers them, with the polarity carefully observed, to the LED-and-resistor combination. The LED lights up brightly, dims visibly over two or three seconds, then fades to dark.

Kipanga: *Sir, it glowed and then it faded.*

Mr. Akilikubwa: *It did. It glowed because the capacitor was driving current through the LED. It faded because the capacitor was running out of charge to drive that current with. By the end of the fade, the voltage across the capacitor had fallen to a value too small to light the LED.*

Kipute: *And the resistor controlled how slowly it faded?*

Mr. Akilikubwa: *It did. A larger resistor would have made the fade slower. A smaller resistor would have made it faster, and probably also burnt out the LED in the first rush of current before any fading could begin. The capacitor itself plays a role too. Replace this 1000-microfarad capacitor with a smaller 100-microfarad one, and the fade is ten times faster. Replace it with a 10 000-microfarad capacitor, and the LED holds its glow for tens of seconds.*

Kipanga: *So the capacitor is a kind of small battery?*

Mr. Akilikubwa: *It behaves like one during the discharge. But unlike a battery, a capacitor stores no chemical energy. It stores only the energy of the electric field between its plates. When the field drains away, so does the energy. A battery, once exhausted, can sometimes be revived by chemistry; a capacitor, once discharged, is simply empty until something charges it again. Later in this chapter we will calculate exactly how long the fade takes for any given combination of capacitor and resistor, using an exponential law and a single number called the time constant. For now you have the picture. A capacitor fills when connected to a battery and empties when connected across a load. The speed of either process depends on the resistance of the path.*

THE PARALLEL-PLATE CAPACITOR

In the previous section we wrote down the defining equation $C = \frac{Q}{V}$, but that equation by itself tells us only what capacitance is. It does not tell us what determines the value of C for any particular device. To predict that value, we have to look inside the capacitor and account for the geometry of the plates and the material that fills the gap between them. We start with the simplest geometry of all: two flat parallel plates separated by a uniform gap.

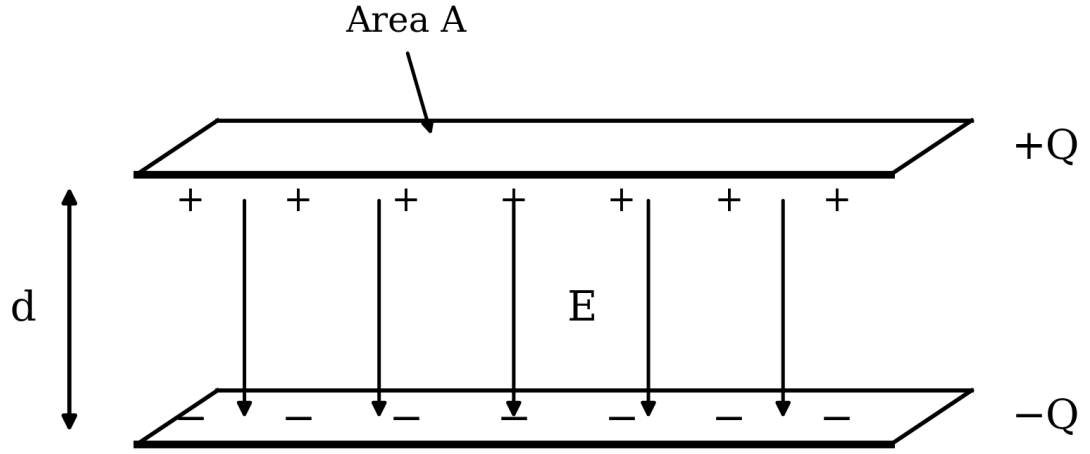


Figure: ***A parallel-plate capacitor.*** *Each plate has area A and the gap between them is d. The top plate carries charge +Q, spread evenly over its area; the bottom plate carries an equal and opposite charge −Q. The field between the plates is uniform and points from the positive plate to the negative one.*

Let each plate have area A, and let the uniform gap between them be d. Suppose the capacitor has been connected to a battery long enough to be fully charged. The top plate carries a positive charge Q distributed evenly over its surface, and the bottom plate carries the same magnitude of charge with the opposite sign, also distributed evenly. The surface charge density on each plate is therefore:

$$\sigma = \frac{Q}{A}$$

In the previous chapter we showed that the electric field between two parallel plates of opposite surface charge has the simple uniform value $E = \frac{\sigma}{\varepsilon_0}$ in the region between the plates and is essentially zero outside, provided the plate dimensions are much larger than the gap d. This proviso is met by almost every real capacitor; the gap is always small compared with the plate dimensions. Substituting our expression for σ:

$$E = \frac{\sigma}{\varepsilon_0} = \frac{Q}{A\varepsilon_0}$$

This is the magnitude of the uniform electric field everywhere in the gap. The potential difference between the two plates is the work done per unit charge to move a test charge from one plate to the other across the gap. Because the field is uniform and the displacement is parallel to it, the potential difference is simply the field strength multiplied by the gap:

$$V = Ed = \frac{Qd}{A\varepsilon_0}$$

Both Q and V depend on how much charge the battery has placed on the plates; the bigger the battery, the bigger both numbers. But their ratio is set entirely by the geometry of the capacitor, with no reference to the battery:

$$\mathbf{C = \frac{Q}{V} = \frac{A\varepsilon_0}{d}}$$

That is the central result for a parallel-plate capacitor. With three symbols on the right-hand side, it tells us three things at once.

1) Large plates make large capacitance

Doubling the plate area A doubles the capacitance, with no change to the voltage rating. This is the most direct way of increasing the capacitance of a parallel-plate device: make the plates bigger. The cost is bulk, since the device itself grows in the same proportion.

2) Close plates make large capacitance

Halving the gap d also doubles the capacitance, this time with no change in the plate area. In some ways this is an even more useful relation, because it works without enlarging the device. But it cannot be pushed without limit: the smaller the gap, the more easily the dielectric between the plates breaks down under the applied field, and the smaller the working voltage at which the capacitor can be operated safely.

3) The constant ε_0 is the property of empty space

No engineering choice changes its value. But if we replace the air or vacuum gap with an insulating material called a dielectric, the effective permittivity rises above ε_0, and the capacitance rises in the same proportion. This is the third method of increasing capacitance, and the most powerful in practice. The whole industry of large-value capacitors is built on the development of dielectrics of high permittivity. We return to this idea in detail in the section on dielectrics later in the chapter.

For now, with the gap full of air or vacuum, the formula stands as $C = \frac{A\varepsilon_0}{d}$, and ε_0 takes its SI value:

$$\varepsilon_0 = 8.854 \times 10^{-12}\,\mathrm{F/m}$$

There is one practical limit on how thin the gap d can be made. At very small separations the electric field $E = \frac{V}{d}$ in the gap becomes so large that the air between the plates ionises, and the capacitor discharges itself across the gap as a brief spark.

The field at which this happens is called the dielectric strength of the medium, and for dry air at ordinary pressure and temperature it has the value $E_{breakdown} \approx 3 \times 10^6\ \mathrm{V/m}$. This is the same field that we met at the end of Chapter 10, when we examined the lightning bolt between a charged cloud base and the ground: lightning is the breakdown of the air column when the field between cloud and ground exceeds this threshold. Inside an air-gap capacitor the same physics sets the maximum voltage the device can withstand for a given gap: $V_{max} = E_{breakdown} \times d$. Push the applied voltage past V_{max} and the capacitor flashes over. We will return to dielectric strengths of other materials in a later section, when we examine what happens when an insulating slab is placed between the plates.

It is also worth pausing on the geometry of the field between the plates. Inside the central region (the part of the gap well away from the plate edges) the field is straight, parallel, and uniform, as we have already seen. The surfaces of constant potential in this region, called the equipotential surfaces, are flat planes parallel to the plates themselves, evenly spaced through the gap; the field lines cross every equipotential at right angles. Near the edges of the plates, however, the field is no longer perfectly uniform. *The lines bulge outward into the space just beyond the plates,* an effect called **edge fringing** or the **fringe field.** The equipotentials curve to follow them. For plates whose linear dimensions are large compared with the gap d, the fringing region is a small fraction of the total field volume, and the uniform-field approximation used in the derivation of $C = \frac{\varepsilon_0 A}{d}$ is excellent. For very small plates with a comparatively large gap, fringing becomes a noticeable correction and more careful methods are required.

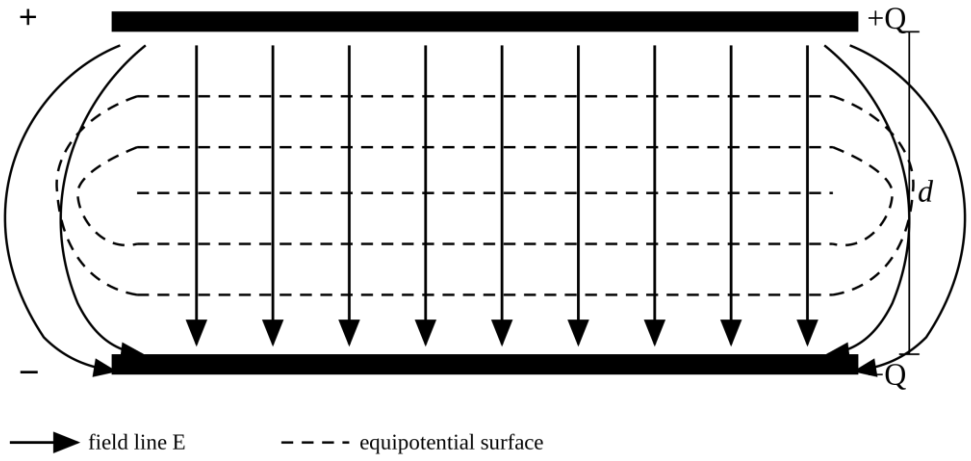


Figure: ***Equipotential surfaces and field lines for a parallel-plate capacitor.*** *The field lines (solid arrows) are straight and evenly spaced in the central region between the plates and curve outward at the edges in the fringing region. The equipotentials (dashed) are flat planes in the central region and curve outward near the edges. Field lines always cross equipotentials at right angles.*

Multi-plate capacitors

A real capacitor can multiply its effective plate area without growing in overall size by stacking many plates and connecting them alternately to the two leads. Suppose we have N plates stacked with a uniform gap d between each adjacent pair, and we wire the odd-numbered plates (1, 3, 5, and so on) to one lead and the even-numbered plates (2, 4, 6, and so on) to the other lead.

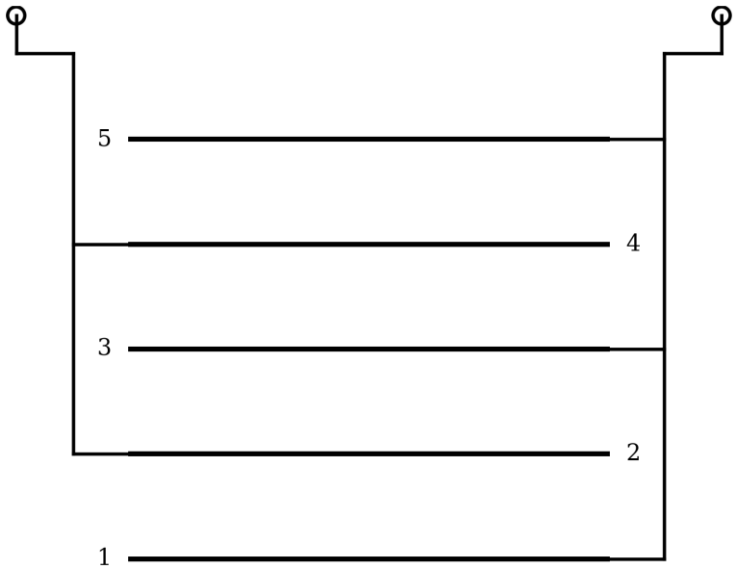


Figure: ***A multi-plate capacitor with five plates.*** *Odd-numbered plates are wired to one terminal, even-numbered plates to the other. Each adjacent pair of plates forms one parallel-plate capacitor, and the four such pairs are all wired in parallel.*

Each adjacent pair of plates now behaves as its own small parallel-plate capacitor with plate area A and gap d. The number of such pairs is one less than the number of plates, since between N plates there are $N - 1$ gaps. (For the five-plate stack drawn here, the four adjacent pairs are 1–2, 2–3, 3–4, and 4–5.) All of these pairs share the same two leads, and so they are wired in parallel. As we will see in the next section, capacitors wired in parallel have their capacitances added directly. The total capacitance of the multi-plate stack is therefore:

$$\mathbf{C = (N - 1)\frac{A\varepsilon_0}{d}}$$

Each additional plate adds one more pair, and so one more effective capacitor, to the parallel sum.

Most **air-variable capacitors** in old radios are built on exactly this principle. A stack of about 10 to 30 thin metal plates is split in two: half of them are fixed in the cabinet, and half are mounted on a rotating shaft. Turning the tuning knob slides the

rotating set in or out of the fixed set, changing the effective overlap area A smoothly between almost zero (plates fully apart) and the full geometrical area (plates fully meshed). The capacitance then varies smoothly between a few picofarads and a few hundred picofarads, which is exactly the range needed to tune a radio across the AM broadcast band.

Five worked examples follow.

BINDER Example 4

A multi-plate mica capacitor is built from 7 thin plates, each of area 50cm^2, stacked with a uniform gap of 0.2mm between consecutive plates. The gaps are filled with mica, whose dielectric constant is $\varepsilon_r = 6$. Find the capacitance of the device.

Solution

With $N = 7$ plates in the stack, the number of adjacent pairs (and so the number of capacitor gaps in parallel) is one less:

$$N - 1 = 7 - 1 = 6$$

Each plate has area $A = 50\text{cm}^2 = 5 \times 10^{-3}\ \text{m}^2$. The gap is $d = 0.2\text{mm} = 2 \times 10^{-4}\ \text{m}$. With mica in the gap, the effective permittivity (ε) is $\varepsilon_r\varepsilon_0$, so the multi-plate formula becomes:

$$C = \frac{(N - 1)\,\varepsilon_r\varepsilon_0 A}{d}$$

Substituting:

$$C = \frac{6 \times 6 \times (8.854 \times 10^{-12}\ \text{F/m}) \times (5 \times 10^{-3}\ \text{m}^2)}{2 \times 10^{-4}\ \text{m}} = 7.97 \times 10^{-9}\ \text{F} = 7.97\text{nF}$$

The capacitance of the seven-plate mica stack is 7.97nF.

Making Sense of the Answer: *A few nanofarads is exactly the order of magnitude one finds for the high-stability mica capacitors used in radio-frequency tuned circuits, where dielectric losses must be tiny and the capacitance must not drift with temperature. The two design moves that pushed our number out of the bare-air picofarad range and into the nanofarad range were the multi-plate stacking (six gaps in parallel rather than one) and the use of a high-ε_r material in the gap (mica multiplies the bare-air value by six). Each of these moves contributes a factor of six in the present example, giving the device a capacitance thirty-six times larger than a single-gap air capacitor of the same plate area would have at the same plate separation.*

Think Like a Physicist: *Multi-plate formulae and dielectric formulae multiply together cleanly, because their underlying physical effects are independent: stacking plates adds capacitor gaps in parallel, while a dielectric in each gap raises that gap's per-pair capacitance. The composite formula $C = \frac{(N-1)\,\varepsilon_r\varepsilon_0 A}{d}$ is best read as the single-pair vacuum capacitance $\frac{\varepsilon_0 A}{d}$, multiplied by $(N - 1)$ for parallel-gap stacking, multiplied by ε_r for the dielectric fill. Identifying these three factors separately on the worksheet makes both the arithmetic and the physical meaning easier to track.*

BINDER Example 5

A parallel-plate capacitor consists of two square metal plates of side 10 cm separated by an air gap of 1 mm. (a) Calculate the capacitance of the device. (b) Find the maximum voltage that can be applied across the plates before air breakdown, given a dielectric strength of $3 \times 10^6\ \text{V/m}$ for dry air. (c) Find the maximum charge the device can hold at this voltage.

Solution

(a) Convert each given length to metres, since the SI unit of ε_0 is built on the metre. The plates are squares of side 10cm, so each has area:

$$A = (0.10\text{m})^2 = 1 \times 10^{-2}\text{m}^2$$

The gap, converted to metres, is:

$$d = 1\text{mm} = 1 \times 10^{-3}\text{m}$$

With $\varepsilon_0 = 8.854 \times 10^{-12}\text{F/m}$, the capacitance is:

$$C = \frac{A\varepsilon_0}{d} = \frac{(1 \times 10^{-2}\text{m}^2)(8.854 \times 10^{-12}\text{F/m})}{1 \times 10^{-3}\text{m}} = 8.854 \times 10^{-11}\text{F} \approx 88.5\text{pF}$$

The capacitance of the device is about 88.5pF.

(b) The maximum voltage that can be applied across the gap before the air ionises is the dielectric strength multiplied by the gap:

$$V_{max} = E_{breakdown} \times d = (3 \times 10^6\ \text{V/m})(1 \times 10^{-3}\ \text{m}) = 3 \times 10^3\ \text{V} = 3000\text{V}$$

The capacitor can safely sustain potential differences up to about 3kV before the air between its plates ionises and the device discharges across the gap in a spark.

(c) The maximum charge the capacitor can hold at this maximum voltage follows from the defining equation of capacitance:

$$Q_{max} = C\,V_{max} = (88.5 \times 10^{-12}\ \text{F})(3000\ \text{V}) = 2.66 \times 10^{-7}\ \text{C} \approx 0.27\mu\text{C}$$

The maximum charge the device can store before discharging through the air is about $0.27\mu\text{C}$.

Making Sense of the Answer: *About ninety picofarads for 10cm by 10cm plates a millimetre apart is exactly the order of magnitude one finds in the variable tuning capacitors of old AM radios. Small, tens to hundreds of picofarads, is the natural*

scale of a hand-built parallel-plate device. To climb into the microfarad range one must take one of three paths: enlarge the plates beyond what a hobbyist can build, shrink the gap to a few micrometres at the price of low working voltage, or fill the gap with a dielectric of high permittivity. The first path leads to bulky equipment, the second to fragile equipment, and the third to the electrolytic capacitor in every power supply on this planet.

Think Like a Physicist: *The constant ε_0 has the awkward value 8.854×10^{-12} F/m, and that minus-twelve hides in every parallel-plate calculation, waiting to be dropped. A useful habit is to write ε_0 in scientific notation at the very top of the working, before substituting anything else, and then carry the powers of ten through the algebra by adding and subtracting their exponents in a clearly visible column. Forgetting a single 10^{-12} gives an answer that is wrong by a factor of a trillion, and on examination scripts this is the single commonest mistake on parallel-plate questions.*

REAL Example 6

In the school prep room, **Kipute** opens a small capacitance meter and clips its leads across an air-gap parallel-plate capacitor whose plates are exposed. She breathes gently on the gap between the plates and watches the meter reading: it shifts only by a few per cent. She then slides a thin sheet of dry paper between the plates (without touching either plate) and watches the reading again: it now changes by a factor of about three. Explain qualitatively why the two interventions produce such different effects.

Solution

Both interventions are doing the same physical thing: altering the dielectric properties of the material that fills the gap. But they do it on very different scales.

Air at ordinary temperature and humidity is almost a perfect vacuum from the point of view of the electric field. Its dielectric constant is $\varepsilon_r \approx 1.0006$ in dry air, and rises only slightly with humidity. When Kipute breathes on the gap, she increases the water-vapour content of the air slightly, which raises ε_r by a fraction of a percent. The capacitor sees a tiny corresponding rise in its capacitance, and the meter responds by a few per cent. Even a wet breath barely lifts ε_r above one.

Paper is a very different beast with greater value of ε_r. Even dry paper has $\varepsilon_r \approx 3$, a value three times the dielectric constant of air. Slipping the paper between the plates replaces most of the air in the gap with a much more easily polarised material; the resulting capacitance is therefore about three times the empty-gap value. The meter responds by a factor of three.

The qualitative lesson is that what counts is not what material is touching the plates, but what material fills the gap between the plates. That is where the field stores its energy and the dielectric does its work. Humid air remains, electrically, almost vacuum. Paper is fundamentally different.

Making Sense of the Answer: *A factor of three in capacitance from a single sheet of paper is the simplest demonstration that the medium between the plates matters as much as the geometry. Breathing on the gap, by contrast, is a stress-test of how close to vacuum ordinary air really is.*

Think Like a Physicist: *Whenever a question contrasts two experimental interventions and asks you to explain why one matters more than the other, the cleanest answer names the underlying parameter that is being varied (here, the dielectric constant ε_r of the material in the gap) and the magnitude of the variation in each case (fractional per-cent for humidity in air, factor of three for paper). The qualitative answer is usually a one-line ratio between the two relevant ε_r values, with the rest of the paragraph spent justifying the ratio.*

REAL Example 7

The variable tuning capacitor of an old AM radio receiver is built from a stack of 21 thin semicircular brass plates of radius 2.5cm, separated by air gaps of 0.5mm. The odd-numbered plates are fixed to the cabinet, and the even-numbered plates are mounted on a tuning shaft so that they rotate in and out of the fixed set. Calculate the maximum capacitance of the device, achieved when the moving plates are fully meshed with the fixed ones.

Solution

With 21 plates in the stack, the number of adjacent pairs is one less:

$$\mathrm{N}-1=21-1=20$$

Each plate is a semicircle of radius $\mathrm{r}=2.5\mathrm{cm}=2.5\times10^{-2}\mathrm{m}$. The area of one semicircle is:

$$\mathrm{A}=\frac{1}{2}\pi\mathrm{r}^2=\frac{1}{2}\pi(2.5\times10^{-2}\mathrm{m})^2=9.817\times10^{-4}\mathrm{m}^2$$

The gap, converted to metres, is $\mathrm{d}=0.5\mathrm{mm}=5\times10^{-4}\mathrm{m}$. The total capacitance of the meshed stack is:

$$\mathrm{C}=(\mathrm{N}-1)\frac{\mathrm{A}\varepsilon_0}{\mathrm{d}}$$

$$\mathrm{C}=20\times\frac{(9.817\times10^{-4}\mathrm{m}^2)(8.854\times10^{-12}\mathrm{F/m})}{5\times10^{-4}\mathrm{m}}=3.48\times10^{-10}\mathrm{F}=348\mathrm{pF}$$

The maximum capacitance of the tuner is about 348pF.

Making Sense of the Answer: *The medium-wave AM broadcast band runs from about 530kHz to 1600kHz. An AM radio tunes across this band by varying the capacitance of a parallel LC circuit; the inductor is a fixed coil of two or three hundred microhenrys, and the variable capacitor has to swing from about 20pF (when the moving plates are fully open) to 350 or 400pF (when fully meshed). Our calculated maximum of 348pF lands exactly where it should. The reader who builds an AM radio at home is therefore building a capacitor of this rough size and shape, with their own hands, and the result of their effort is the ability to select a station out of a few dozen contenders by rotating a single knob.*

Think Like a Physicist: *A multi-plate capacitor's capacitance is the per-pair capacitance multiplied by the number of pairs, not by the number of plates. With 21 plates the number of pairs is 20, with the simple rule that the number of gaps is one less than the number of plates. Forgetting this off-by-one is the single commonest mistake on multi-plate questions. A student who recognises a multi-plate problem on sight should write the quantity (N − 1) on the working sheet before doing anything else, just as a student who recognises a circular wire problem on sight should write down $\pi d^2/4$ before doing anything else.*

HOT Example 8

A parallel-plate capacitor consists of two circular metal plates of radius 5cm separated by an air gap of 2mm. The capacitor is connected to a 50V supply and allowed to charge fully. The supply is then disconnected. After the supply has been removed, the gap between the plates is reduced to 0.5mm by a mechanical adjustment. Find (a) the original capacitance and the original charge on each plate; (b) the new capacitance after the gap has been reduced; (c) the new voltage between the plates.

Solution

Before any algebra, write down the central physical fact about this problem: the supply has been disconnected before the geometry change. With the supply gone, there is no path for charge to flow on or off the plates, and so the total charge Q on each plate is held fixed throughout the gap change. The capacitance C changes because the geometry changes; the voltage V then adjusts to keep the relation $\text{Q} = \text{CV}$ consistent.

(a) The plate area is:

$$\text{A} = \pi\text{r}^2 = \pi(5 \times 10^{-2}\text{m})^2 = 7.854 \times 10^{-3}\text{m}^2$$

With the original gap $\text{d}_{\text{old}} = 2\text{mm} = 2 \times 10^{-3}\text{m}$, the original capacitance is:

$$\text{C}_{\text{old}} = \frac{\text{A}\varepsilon_0}{\text{d}_{\text{old}}}$$

$$\text{C}_{\text{old}} = \frac{(7.854 \times 10^{-3}\text{m}^2)(8.854 \times 10^{-12}\text{F/m})}{2 \times 10^{-3}\text{m}} = 3.48 \times 10^{-11}\text{F} = 34.8\text{pF}$$

The original charge on each plate is:

$$\text{Q} = \text{C}_{\text{old}}\text{V}_{\text{old}} = (34.8 \times 10^{-12}\text{F})(50\text{V}) = 1.74 \times 10^{-9}\text{C} = 1.74\text{nC}$$

(b) The new gap is $\text{d}_{\text{new}} = 0.5\text{mm} = 5 \times 10^{-4}\text{m}$, which is one quarter of the original gap. Since the capacitance is inversely proportional to the gap (with everything else held fixed), reducing the gap by a factor of four multiplies the capacitance by four:

$$\text{C}_{\text{new}} = 4\text{C}_{\text{old}} = 4 \times 34.8\text{pF} = 139.2\text{pF}$$

(c) The charge Q is held fixed (the supply was disconnected before the gap was changed). The new voltage is therefore:

$$\text{V}_{\text{new}} = \frac{\text{Q}}{\text{C}_{\text{new}}} = \frac{1.74 \times 10^{-9}\text{C}}{139.2 \times 10^{-12}\text{F}} = 12.5\text{V}$$

After the gap is reduced to 0.5mm, the capacitance rises to about 139pF and the voltage between the plates falls to about 12.5V, while the charge of 1.74nC remains undisturbed on each plate.

Making Sense of the Answer: *Reducing the gap by a factor of four has increased the capacitance by a factor of four (because C is inversely proportional to d), and with the charge held fixed by the disconnected supply, the voltage must drop by a factor of four to keep the relation Q = CV satisfied. Fifty volts divided by four is twelve and a half volts, which is what we obtained. The total charge of 1.74nC sat undisturbed on the plates throughout the gap change; the geometry change only altered the ease with which that charge could be held at a given voltage. The problem is, at heart, the relation Q = CV restated three times in three different roles.*

Think Like a Physicist: *Every "what happens after the supply is disconnected" question is asking which quantities are held fixed by the disconnection (always: the charge on each plate) and which quantities are free to change (always: the voltage, if the capacitance changes; and the capacitance, if the geometry changes). The first move on any such question is to write the words "**Q is held fixed**" at the top of the working before any algebra begins. The opposite pattern, in which the supply stays connected throughout the change, holds the voltage fixed instead and lets the charge change. Knowing which of the three quantities (Q, V, C) is the constant of the problem decides the entire solution; getting it wrong sends every subsequent line of working in the wrong direction.*

We now know how to compute the capacitance of a single parallel-plate device of any geometry, and the capacitance of a stack of such devices acting as one. The next section asks the inevitable follow-up question: *what happens when several finished capacitors are wired together in a circuit?*

COMBINATIONS OF CAPACITORS

In a perfect engineering world, every capacitor on the market would be available in exactly the value a designer wants. In the real world, capacitors come in standard sizes (1μF, 2.2μF, 4.7μF, 10μF, and so on), and the designer who needs, say, 4.5μF or 13.6μF must combine two or more standard parts to obtain the value required.

There are two ways to combine capacitors in a circuit. They can be wired in parallel, in which case both plates of each capacitor are connected across the same pair of nodes; or they can be wired in series, in which case the capacitors form a single chain between two nodes. The two arrangements produce very different total capacitances, and the way they share charge and voltage is different again. We deal with each in turn, and then with the related question of what happens when two capacitors that have been charged separately are afterwards connected together, the so-called charge-sharing problem.

Capacitors in parallel

When two or more capacitors are wired in parallel, both plates of each capacitor are connected across the same pair of nodes in the circuit. The same pair of wires links every capacitor's top plate, and another pair links every capacitor's bottom plate. Every capacitor in the combination therefore experiences the same potential difference V across its terminals: the voltage of the supply.

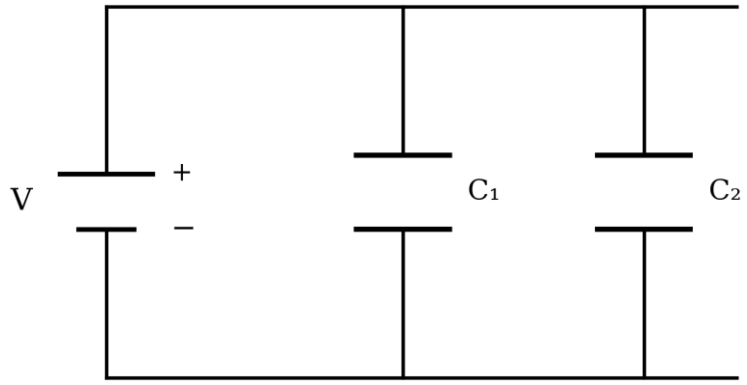


Figure: ***Two capacitors in parallel across a battery.*** *Each capacitor sees the full battery voltage V across its plates, but each stores its own charge in proportion to its own capacitance.*

The charge each capacitor accumulates is proportional to its own capacitance:

$$Q_1 = C_1V, \quad Q_2 = C_2V, \quad Q_3 = C_3V, \quad \ldots$$

The total charge stored across the combination is the sum of the charges on each individual capacitor:

$$Q_{\text{total}} = Q_1 + Q_2 + Q_3 + \cdots = (C_1 + C_2 + C_3 + \cdots)V$$

Defining the **equivalent capacitance $\mathbf{C_{eq}}$** as the single capacitance that, when connected to the same supply, would store the same total charge:

$$Q_{\text{total}} = C_{\text{eq}}V$$

Comparing the two expressions for Q_{total}:

$$C_{\text{eq}}V = (C_1 + C_2 + C_3 + \cdots)V$$

$$\mathbf{C_{eq} = C_1 + C_2 + C_3 + \cdots}$$

In parallel, capacitances add directly. This is the opposite of resistors, where parallel resistances combine reciprocally. The asymmetry is not a coincidence: *capacitors store charge while resistors pass it.* More capacitors in parallel provide more charge storage, in the same way that more buckets standing side by side hold more water; more resistors in parallel provide less resistance, in the same way that more lanes on a road create less obstruction to traffic flow.

Capacitors in series

When two or more capacitors are wired in series, they form a single chain between two nodes. The right plate of one capacitor is connected to the left plate of the next, and so on along the chain, with no other path for charge to enter or leave the middle wires.

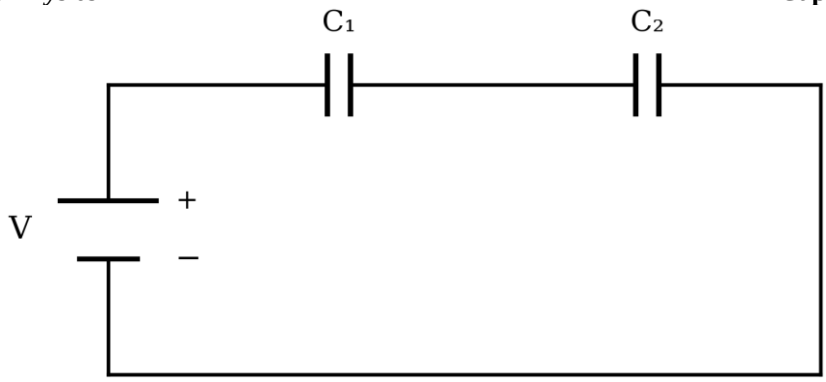


Figure: ***Two capacitors in series across a battery.*** *The middle wire between the two capacitors is isolated from the rest of the circuit, and the charge on every plate in the chain has the same magnitude.*

This isolation of the middle wires is the key to understanding the series combination. When the supply pushes charge $+Q$ onto the leftmost plate of the chain, the right plate of that same capacitor must carry $-Q$. The bottom-plate charge cannot come from the supply, because the supply is not connected to the middle wires. Instead, electrons from the next capacitor's left plate are drawn through the middle wire to settle on the first capacitor's right plate. The next capacitor's left plate, having lost those electrons, becomes $+Q$. The same argument repeats along the entire chain. Every capacitor in the series chain therefore carries the same magnitude of charge Q on its plates. This is the central fact about series-connected capacitors.

The voltage across each individual capacitor in the chain is set by its own capacitance:

$$V_1 = \frac{Q}{C_1}, \quad V_2 = \frac{Q}{C_2}, \quad V_3 = \frac{Q}{C_3}, \quad \ldots$$

The total voltage across the chain is the sum of these individual voltages:

$$V = V_1 + V_2 + V_3 + \cdots = Q\left(\frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} + \cdots\right)$$

Defining the **equivalent capacitance C_{eq}** as the single capacitor that would carry the same charge Q at the same total voltage V:

$$V = \frac{Q}{C_{eq}}$$

Comparing the two expressions for V:

$$\frac{Q}{C_{eq}} = Q\left(\frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} + \cdots\right)$$

$$\boldsymbol{\frac{1}{C_{eq}} = \frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} + \cdots}$$

In series, the reciprocals of the capacitances add, never the capacitances themselves.

For the common special case of just two capacitors in series, the reciprocal-sum formula rearranges into a more convenient form. Starting from:

$$\frac{1}{C_{eq}} = \frac{1}{C_1} + \frac{1}{C_2} = \frac{C_1 + C_2}{C_1 C_2}$$

Inverting both sides:

$$\boldsymbol{C_{eq} = \frac{C_1 C_2}{C_1 + C_2}}$$

This is the "product over sum" rule, valid for two capacitors in series and **only** two.

It is worth recording the closed form for three capacitors in series as well. Starting from:

$$\frac{1}{C_{eq}} = \frac{1}{C_1} + \frac{1}{C_2} + \frac{1}{C_3} = \frac{C_2C_3 + C_1C_3 + C_1C_2}{C_1C_2C_3}$$

Inverting both sides:

$$\boldsymbol{C_{eq} = \frac{C_1C_2C_3}{C_1C_2 + C_1C_3 + C_2C_3}}$$

For four or more capacitors the closed form becomes unwieldy and the reciprocal sum must be evaluated directly, or the capacitors combined two at a time.

Before leaving the series formula, one observation is worth pinning down.

In a series chain, every capacitor carries the same charge Q (charge conservation through the chain), so the voltage across each capacitor is $V_i = \frac{Q}{C_i}$. The capacitor with the smallest C therefore carries the largest V. *Series capacitors divide the supply voltage in inverse proportion to their capacitances.*

By the same logic, in a parallel combination every capacitor sees the same V across its plates, so the charge on each is $Q_i = C_i V$. The capacitor with the largest C therefore carries the largest Q. *Parallel capacitors divide the supplied charge in direct proportion to their capacitances.*

These two observations are the engine of every exam-style potential-divider problem and are worth remembering in this compact form: ***in series, smallest C takes the largest V; in parallel, largest C takes the largest Q.***

Charge sharing

A peculiar but instructive scenario arises when two capacitors are charged separately and then connected together. Suppose a capacitor C_1 has been charged to a voltage V_1 by one source, and a second capacitor C_2 has been charged separately to a different voltage V_2 by another source (or by the same source set to a different output). Both sources are then disconnected, and the two capacitors are wired together so that their positive plates are joined by one wire and their negative plates by another. The two capacitors are now in parallel, but each is initially at a different voltage.

Charge cannot remain in this disequilibrium. Electrons flow from the higher-voltage capacitor, through the connecting wire, and onto the lower-voltage one until both come to a common voltage. *What is that common voltage?*

The key idea is conservation of charge. Before the connection, the total charge held on the (originally) positive plates is:

$$Q_{total} = Q_1 + Q_2 = C_1V_1 + C_2V_2$$

After the connection, charge has redistributed between the two capacitors, but the total charge on the connected side is unchanged: nothing has been added to the system, and nothing has left it. The two capacitors are now in parallel with a common voltage V_{final}:

$$Q_{total} = (C_1 + C_2)V_{final}$$

Solving for the final voltage:

$$C_1V_1 + C_2V_2 = (C_1 + C_2)V_{final}$$

$$\mathbf{V_{final} = \frac{C_1V_1 + C_2V_2}{C_1 + C_2}}$$

This is a weighted average of V_1 and V_2, with the weights being the two capacitances. The result always lies between the two initial voltages, and is closer to whichever capacitor is the larger of the two: the bigger capacitor dominates the final voltage, in the same way that a large reservoir dominates the level of a small tank joined to it by a pipe.

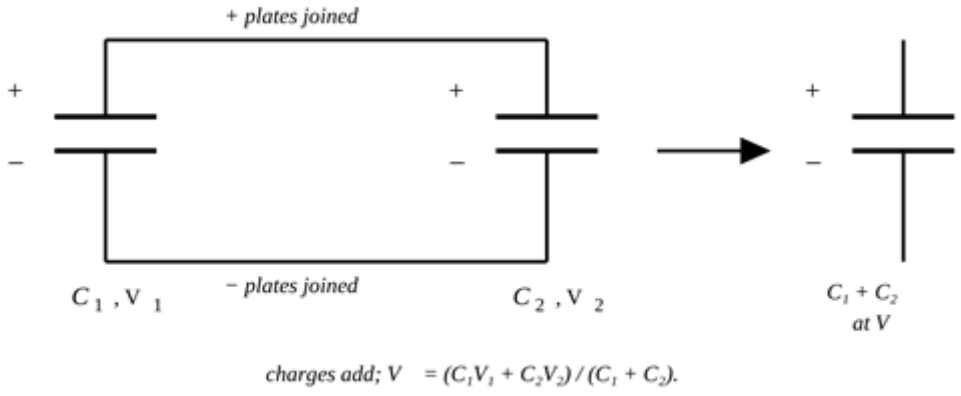


Figure: ***Charge sharing.*** *Like plates joined: the positive plate of each capacitor is wired to the positive plate of the other, so charges of the same sign add.*

A warning: The stated formula assumes the two capacitors are connected positive plate to positive plate, and negative to negative. If they are instead connected with opposite polarity (positive to negative), the charges on the joined plates partly cancel, and the formula takes the modified form:

$$\mathbf{V_{final} = \frac{|C_1V_1 - C_2V_2|}{C_1 + C_2}}$$

The polarity of the connection therefore matters.

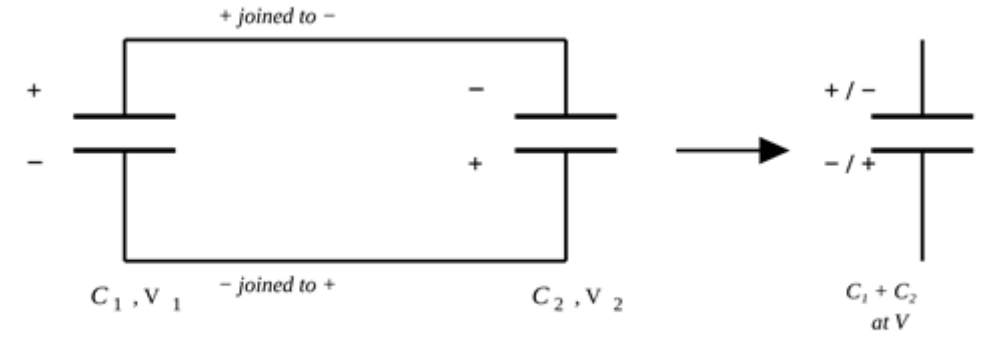


Figure: ***Charge sharing.*** *Unlike plates joined: the positive plate of one capacitor is wired to the negative plate of the other, so charges of opposite sign partially cancel.*

Five worked examples now follow.

BINDER Example 9

A 2μF capacitor and a 3μF capacitor are wired in parallel across a 10V battery. Find (a) the equivalent capacitance of the combination; (b) the charge stored on each capacitor; (c) the total charge supplied by the battery.

Solution

(a) Parallel capacitances add directly:

$$C_{eq} = C_1 + C_2 = 2\mu F + 3\mu F = 5\mu F$$

(b) In a parallel combination, every capacitor sees the full battery voltage $V = 10V$. The charge on each is:

$$Q_1 = C_1 V = (2\mu F)(10V) = 20\mu C$$
$$Q_2 = C_2 V = (3\mu F)(10V) = 30\mu C$$

(c) The total charge supplied by the battery is the sum of the charges stored on the two capacitors:

$$Q_{total} = Q_1 + Q_2 = 20\mu C + 30\mu C = 50\mu C$$

Alternatively, applying the parallel formula directly: $Q_{total} = C_{eq} V = (5\mu F)(10V) = 50\mu C$, in agreement.

Making Sense of the Answer: *The arithmetic is undramatic, but two features are worth noting. First, every capacitor in the parallel combination carries the same voltage, in the same way that every lamp wired in parallel across a supply glows at the supply voltage. Second, the larger capacitor stores more charge, in proportion to its capacitance: the 3μF capacitor stores 30μC, the 2μF stores 20μC, and the ratio of charges is the ratio of capacitances.*

Think Like a Physicist: *A good cross-check on any parallel-capacitor question is to compute the total charge in two ways: once as the sum of the individual charges (which uses every individual capacitance), and once as (equivalent capacitance)(voltage) (which uses only the equivalent). The two answers must agree. If they do not, an arithmetic mistake is somewhere in the working, and the cross-check has caught it before any mark has been lost. This double-bookkeeping habit takes ten extra seconds and saves marks that would otherwise leak out of the script unnoticed.*

BINDER Example 10

A 4μF capacitor and a 6μF capacitor are wired in series across a 10V battery. Find (a) the equivalent capacitance of the combination; (b) the charge stored on each capacitor; (c) the voltage across each capacitor.

Solution

(a) Using the product-over-sum formula for two capacitors in series:

$$C_{eq} = \frac{C_1 C_2}{C_1 + C_2} = \frac{(4\mu F)(6\mu F)}{4\mu F + 6\mu F} = 2.4\mu F$$

(b) In a series combination, every capacitor in the chain carries the same charge, equal to the charge that the equivalent capacitor would carry at the supply voltage:

$$Q = C_{eq} V = (2.4\mu F)(10V) = 24\mu C$$

Both capacitors carry the same charge $Q = 24\mu C$ on their plates.

(c) The voltage across each capacitor is set by its own capacitance:

$$V_1 = \frac{Q}{C_1} = \frac{24\mu C}{4\mu F} = 6V$$
$$V_2 = \frac{Q}{C_2} = \frac{24\mu C}{6\mu F} = 4V$$

As a cross-check, $V_1 + V_2 = 6V + 4V = 10V$, which is the supply voltage. The chain has divided the supply voltage between its two members in inverse proportion to their capacitances.

Making Sense of the Answer: *Three lessons are buried in this short calculation. First, the equivalent capacitance of a series combination is always smaller than the smallest individual capacitance (2.4μF here is smaller than 4μF, and very much smaller than 6μF). Series wiring weakens a capacitor; parallel wiring strengthens it. Second, the voltage is shared between the two capacitors in inverse proportion to their capacitances: the smaller capacitor holds more of the voltage (6V across 4μF, 4V across 6μF), because for the same charge, a smaller capacitance needs a larger voltage. Third, the two voltages must add to the supply voltage, which is the most reliable cross-check available on any examination problem.*

Think Like a Physicist: *The two-capacitor product-over-sum rule ($C_1C_2/(C_1 + C_2)$) is fast and reliable. Use it whenever exactly two capacitors are in series. For three or more, do not look for a three-capacitor shortcut; there is no clean one. Either compute the reciprocal sum in full or combine the capacitors two at a time, applying the two-capacitor formula iteratively. Trying to invent a three-capacitor shortcut on the fly during an examination is a guaranteed source of arithmetic mistakes.*

REAL Example 11

In the circuit on a hobbyist's bench, three capacitors are wired together as follows: a $4\mu F$ capacitor C_1 sits in parallel with an $8\mu F$ capacitor C_2, and this parallel pair is wired in series with a $6\mu F$ capacitor C_3. The whole arrangement is then connected across a $12V$ supply. Find (a) the equivalent capacitance of the entire circuit; (b) the charge on each capacitor; (c) the voltage across each capacitor.

Solution

(a) First combine C_1 and C_2, which are in parallel:

$$C_{12} = C_1 + C_2 = 4\mu F + 8\mu F = 12\mu F$$

Now combine C_{12} in series with C_3 using the product-over-sum rule:

$$C_{eq} = \frac{C_{12}C_3}{C_{12} + C_3} = \frac{(12\mu F)(6\mu F)}{12\mu F + 6\mu F} = 4\mu F$$

(b) In the series part of the circuit, the same charge passes through C_{12} and C_3. This charge is equal to the charge on the equivalent capacitor at the supply voltage:

$$Q_{series} = C_{eq}V = (4\mu F)(12V) = 48\mu C$$

The charge on C_3 is therefore $Q_3 = 48\mu C$. The same magnitude of charge sits across the parallel pair, but it splits between C_1 and C_2 in proportion to their individual capacitances. The voltage across the parallel pair is found first:

$$V_{12} = \frac{Q_{series}}{C_{12}} = \frac{48\mu C}{12\mu F} = 4V$$

Now the charges on the two parallel capacitors:

$$Q_1 = C_1V_{12} = (4\mu F)(4V) = 16\mu C$$
$$Q_2 = C_2V_{12} = (8\mu F)(4V) = 32\mu C$$

Cross-check: $Q_1 + Q_2 = 16\mu C + 32\mu C = 48\mu C$, which equals the series charge as it must.

(c) The voltage across C_3 is:

$$V_3 = \frac{Q_{series}}{C_3} = \frac{48\mu C}{6\mu F} = 8V$$

Both C_1 and C_2 carry the same voltage as the parallel pair, $V_{12} = 4V$.

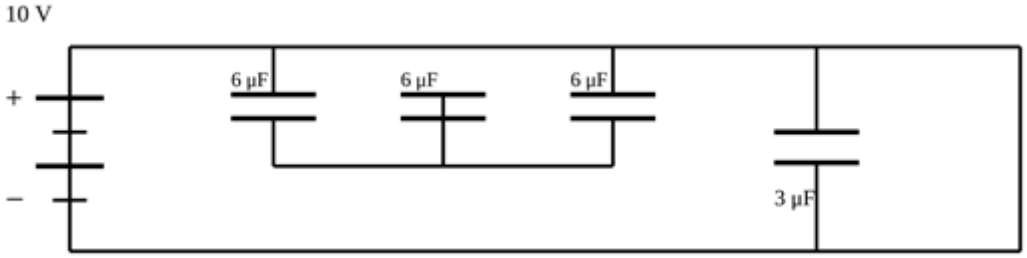
Cross-check on the series voltage rule: $V_{12} + V_3 = 4V + 8V = 12V$, the supply voltage.

Making Sense of the Answer: *Three quick lessons appear in this answer. First, the equivalent capacitance of the whole circuit (4μF) is exactly equal to the individual capacitance of C_1. This is a coincidence of the chosen numbers; do not read any general meaning into it. Second, the smaller member of the series part (the 6μF capacitor C_3) takes the larger share of the supply voltage (8V out of 12V), in exact accordance with the rule that voltage in series divides in inverse proportion to capacitance. Third, the parallel pair shares its 48μC in the ratio of its two capacitances, with the larger 8μF capacitor taking the larger 32μC share. Every cross-check in this problem passed; if any one of them had failed, an arithmetic mistake would have been hiding in the working.*

Think Like a Physicist: *A compound capacitor problem is a sequence of two-capacitor problems, each handled with the appropriate parallel or series rule. The discipline is to combine in stages: identify the innermost group, reduce it to a single equivalent, then move outward. At each stage, write down the new combined capacitance and carry it forward as a single labelled quantity (here, C_{12}) rather than substituting its numerical value into the next equation. Carrying symbolic labels prevents arithmetic errors from accumulating through the stages, and it also makes it easy to see which individual capacitor is responsible for which physical effect at the end.*

HOT Example 12

Use the following figure to find: (a) the equivalent capacitance of the whole combination; (b) the charge stored on the 3 μF capacitor; and (c) the charge on each of the 6 μF capacitors.



Solution

(a) The three 6μF capacitors are in series. Using the reciprocal-sum formula:

$$\frac{1}{C_{series}} = \frac{1}{6\mu F} + \frac{1}{6\mu F} + \frac{1}{6\mu F} = \frac{3}{6\mu F} = \frac{1}{2\mu F}$$

So $C_{series} = 2\mu F$.

The series stack is now in parallel with the 3μF capacitor, and parallel capacitances add directly:

$$C_{eq} = C_{series} + 3\mu F = 2\mu F + 3\mu F = 5\mu F$$

(b) The 3μF capacitor sees the full supply voltage of 10V, so its charge is:

$$Q_{3\,\mu F} = CV = (3\mu F)(10V) = 30\mu C$$

(c) The series stack also sees 10 V across its terminals as a whole, so the charge on the equivalent 2 μF capacitor is:

$$Q_{series} = (2\mu F)(10\,V) = 20\mu C$$

Every capacitor inside the series chain carries this same charge, since the chain is a single closed loop with no branching: each of the three 6μF capacitors carries 20μC on its plates.

Making Sense of the Answer: *Three identical 6μF capacitors in series behave as a single 2μF capacitor: the series chain divides each individual capacitance by the number of capacitors in the chain when those capacitors are equal. Adding a parallel 3μF lifts the combination to 5μF. The pattern recurs constantly in practical circuit design: series stacking lets a designer build a high-voltage capacitor from a chain of lower-voltage units, at the cost of a smaller total capacitance, and a parallel branch can then be used to top up the total to a desired value.*

Think Like a Physicist: *A compound capacitor network is always disentangled by combining innermost groups first. Identify any sub-network whose capacitors are all in pure series or all in pure parallel, replace that sub-network with its single equivalent, redraw the circuit, and repeat. Within three or four redraws even an intimidating-looking network reduces to a single equivalent capacitor across the supply. The discipline of redrawing at every stage is the only thing that keeps the bookkeeping clean on harder problems.*

HOT Example 13

A 4μF capacitor is charged to a voltage of 12V by one source. Separately, a 6μF capacitor is charged to a voltage of 18V by another source. Both sources are then disconnected. The two capacitors are wired together so that their positive plates are joined by one wire and their negative plates by another. Find (a) the final common voltage across the two capacitors; (b) the new charge on each capacitor; (c) the magnitude of charge that flowed from one capacitor to the other during the redistribution.

Solution

(a) Before the connection, the charge on each capacitor is:

$$Q_{1,initial} = C_1V_1 = (4\mu F)(12V) = 48\mu C$$
$$Q_{2,initial} = C_2V_2 = (6\mu F)(18V) = 108\mu C$$

After the connection, charge has redistributed but the total charge on the connected (positive) side is unchanged:

$$Q_{total} = Q_{1,initial} + Q_{2,initial} = 48\mu C + 108\mu C = 156\mu C$$

The two capacitors are now in **parallel**; their equivalent capacitance is:

$$C_{eq} = C_1 + C_2 = 4\mu F + 6\mu F = 10\mu F$$

The common final voltage on the two capacitors is:

$$V_{final} = \frac{Q_{total}}{C_{eq}} = \frac{156\mu C}{10\mu F} = 15.6V$$

(b) The new charge on each capacitor at this common voltage is:

$$Q_{1,new} = C_1V_{final} = (4\mu F)(15.6V) = 62.4\mu C$$
$$Q_{2,new} = C_2V_{final} = (6\mu F)(15.6V) = 93.6\mu C$$

Cross-check on charge conservation: $Q_{1,new} + Q_{2,new} = 62.4\mu C + 93.6\mu C = 156\mu C$, which equals Q_{total} as it must.

(c) The charge that flowed from C_2 to C_1 during the redistribution is the difference between the initial and final charges on C_1:

$$\Delta Q = Q_{1,new} - Q_{1,initial} = 62.4\mu C - 48\mu C = 14.4\mu C$$

An identical magnitude of charge ($14.4\mu C$) flowed off C_2, since $Q_{2,initial} - Q_{2,new} = 108\mu C - 93.6\mu C = 14.4\mu C$, in agreement with charge conservation.

Making Sense of the Answer: *The final voltage of 15.6V sits between the two initial voltages of 12V and 18V, as expected: charge can only redistribute, never appear from nothing or vanish into nothing. The final voltage is closer to 18V than to 12V (15.6V is about 60% of the way from 12V to 18V), and the reason is that the larger capacitor was the one that started at 18V: the 6μF capacitor carries more weight in the weighted average that determines V_{final}. The 4μF capacitor's voltage rose by 3.6V, gaining 14.4μC of charge; the 6μF capacitor's voltage fell by 2.4V, losing the same 14.4μC. The smaller capacitor's voltage rose more sharply because the same transferred charge has a larger effect on a smaller capacitance.*

Think Like a Physicist: *Every charge-sharing problem reduces to a single principle: conservation of charge between the moment before the connection and the moment after. Begin every such problem by writing down the total charge on the connected side before the connection. Then equate it to (sum of capacitances)(final voltage) after the connection. The whole problem is then one equation in one unknown. A common trap, set by question setters with mischievous intentions, is to specify the connection as positive plate to negative plate, which makes the initial charges partly cancel. Always read the question carefully on this point: the polarity of the connection is part of the physics, not a stylistic detail. For the reversed-polarity case mentioned in the derivation, the total charge after the connection is:*

$$Q_{total,reversed} = |C_1V_1 - C_2V_2|$$

and the final voltage is correspondingly smaller, sometimes dramatically so.

Combinations of capacitors are now in hand. The next section asks *what is being stored when a capacitor holds a charge, and how much of it.*

ENERGY STORED IN A CAPACITOR

A capacitor that holds a charge Q at a voltage V holds something else as well: **energy**. The work done by the source to push that charge onto the plates does not vanish at the end of charging; it is stored in the electric field between the plates, ready to be released the moment a path is provided. The camera flash, the defibrillator paddle, the spark plug of a petrol engine, and the smoothing capacitor of every laptop charger: all of these are devices that exist specifically to take energy slowly from a low-voltage supply and release it suddenly when the user demands it.

To find how much energy a capacitor stores, imagine the capacitor partway through being charged. At some intermediate instant the plates carry a charge q (less than the final charge Q), and the voltage at this instant is:

$$v = \frac{q}{C}$$

The source now pushes a tiny extra charge dq onto the positive plate, against the voltage v that the already-accumulated charge is producing. The work done by the source in this tiny step is the voltage times the charge moved:

$$dW = vdq = \frac{q}{C}dq$$

The total work done to charge the capacitor from empty ($q = 0$) to fully charged ($q = Q$) is the integral of this expression from 0 to Q:

$$W = \int_0^Q \frac{q}{C}dq = \left[\frac{q^2}{2C}\right]_0^Q = \frac{Q^2}{2C}$$

All of this work has been stored in the capacitor as the energy of the electric field between its plates. We call this stored energy U:

$$\mathbf{U = \frac{Q^2}{2C}}$$

The same result can be obtained by a picture. The voltage during the charging process grows in direct proportion to the charge already on the plates: $v = q/C$ is a straight line through the origin. If we plot v against q, as charge climbs from 0 to Q and voltage climbs from 0 to V, the work done is the area under this straight line. That area is a triangle of base Q and height V.

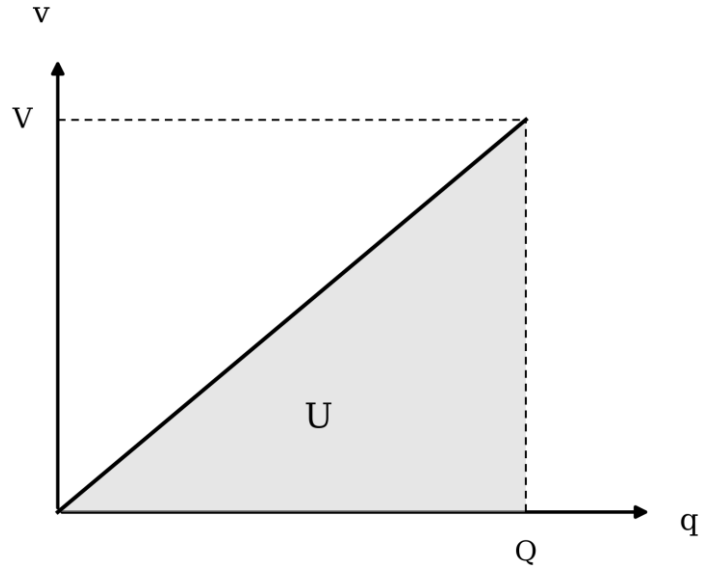


Figure: ***During charging, the voltage v grows linearly with the charge q.*** *The shaded triangle has area* $\frac{1}{2}QV$*, and this area is the energy stored in the capacitor when it is fully charged.*

The area of a triangle of base Q and height V is:

$$\mathbf{U = \frac{1}{2}QV}$$

Two further forms follow at once from the relation $Q = CV$. Substituting $Q = CV$ in the formula above:

$$\mathbf{U = \frac{1}{2}QV = \frac{1}{2}(CV)V = \frac{1}{2}CV^2}$$

Or, substituting $V = Q/C$ in the same expression:

$$U = \frac{1}{2}QV = \frac{1}{2}Q \times \frac{Q}{C} = \frac{Q^2}{2C}$$

The energy stored in a capacitor can therefore be written in three equivalent forms, depending on which two of the three quantities (Q, V, C) are most naturally available in the problem at hand:

$$\mathbf{U = \frac{1}{2}QV = \frac{1}{2}CV^2 = \frac{Q^2}{2C}}$$

The first form is the most general; it relies on knowing both Q and V. The second is most useful when the capacitor is connected to a fixed voltage supply, in which case V is given directly. The third is most useful when the supply has been disconnected and the charge Q is held fixed by isolation. All three forms describe the same physical quantity; substituting $Q = CV$ takes any one of them to either of the other two.

It is illuminating to ask not how much energy a particular capacitor stores, but how much energy is stored per unit volume of the field between its plates. For a parallel-plate capacitor of plate area A, gap d, and field $E = V/d$ between the plates, the energy is:

$$U = \frac{1}{2}CV^2 = \frac{1}{2}\left(\frac{A\varepsilon_0}{d}\right) \times (Ed)^2 = \frac{1}{2}\varepsilon_0 E^2 (Ad)$$

The product Ad is the volume of space between the plates, in which the field E resides. Dividing the energy by this volume gives *the energy stored per unit volume of the field*, called the **energy density** and written **u**:

$$\mathbf{u = \frac{\text{Total energy stored}}{\text{Volume of electric field}} = \frac{U}{Ad} = \frac{1}{2}\varepsilon_0 E^2}$$

This is a general result, not restricted to parallel-plate capacitors. Wherever an electric field E exists, the energy stored in that field, per unit volume of space, is $\frac{1}{2}\varepsilon_0 E^2$. This formula gives a meaning to the field itself as a storehouse of energy, independent of whatever distribution of charges produced it.

Force between the plates

Two oppositely-charged plates attract each other. *How strong is the attraction?* The cleanest way to find out is to use the energy formula. Consider a parallel-plate capacitor with charge Q on each plate, isolated from its supply so that no charge can come or go. The stored energy at gap d is:

$$U(d) = \frac{Q^2}{2C} = \frac{Q^2 d}{2A\varepsilon_0}$$

Now imagine pulling the plates apart by a tiny extra distance dd, doing positive mechanical work against the attractive force. The work done in this small move is $dW = F\,dd$, and since no charge can enter or leave the isolated capacitor, this work goes entirely into increasing the stored energy:

$$F\,dd = dU = \frac{Q^2}{2A\varepsilon_0}\,dd$$

Dividing both sides by dd, the magnitude of the attractive force between the plates is:

$$\mathbf{F = \frac{Q^2}{2A\varepsilon_0}}$$

This can be rewritten in terms of the field E between the plates. Using $Q = A\varepsilon_0 E$ (which follows from $Q = CV$ with $C = A\varepsilon_0/d$ and $V = Ed$):

$$\mathbf{F = \frac{A^2\varepsilon_0^2 E^2}{2A\varepsilon_0} = \frac{1}{2}\varepsilon_0 E^2 A}$$

The force per unit area is therefore $\mathbf{\frac{F}{A} = \frac{1}{2}\varepsilon_0 E^2}$, which is exactly the energy density u of the field between the plates. Force per area equals energy per volume; a deep coincidence rooted in the principle of virtual work, and one that the reader will meet again whenever a system stores energy in a field that can be deformed.

The factor of $\frac{1}{2}$ in $F = \frac{Q^2}{2\,\varepsilon_0\,A}$ has a simple physical origin worth pausing over. The total field in the gap of a parallel-plate capacitor has magnitude $\frac{\sigma}{\varepsilon_0}$, where $\sigma = \frac{Q}{A}$ is the surface charge density. But this total field is the sum of equal contributions from the two plates: each plate alone produces a field of magnitude $\frac{\sigma}{2\varepsilon_0}$, and the two contributions add in the gap (where they point in the same direction) to give the total $\frac{\sigma}{\varepsilon_0}$.

When we ask for the force on one plate, we must use only the field produced by the other plate, since a plate cannot exert a force on itself. The field acting on either plate is therefore $\frac{\sigma}{2\varepsilon_0}$, and the force on the plate of charge Q is $F = Q \times \frac{\sigma}{2\varepsilon_0} = \mathbf{\frac{Q^2}{2A\varepsilon_0}}$. Hence the factor of $\frac{1}{2}$ in the answer is the geometric record of the fact that exactly half of the gap field is produced by the very plate one is computing the force on. We can cross-check the energy-method result against this direct one: both routes give the same F, and the agreement is the consistency between virtual-work bookkeeping and direct field-charge calculation.

Seven worked examples follow. None of them is harder than it looks.

BINDER Example 14

A capacitor of capacitance $5\mu F$ is charged to a potential difference of $12V$. Calculate the energy stored in three different ways: using $U = \frac{1}{2}CV^2$; using $U = \frac{1}{2}QV$; and using $U = \frac{Q^2}{2C}$. Show that all three forms yield the same numerical answer.

Solution

First compute the charge on the plates, which is needed for two of the three formulae:

$$Q = CV = (5\mu F)(12V) = 60\mu C$$

Now compute the energy three ways. Using the second form:

$$U = \frac{1}{2}CV^2 = \frac{1}{2}(5\mu F)(12V)^2 = \frac{1}{2}(5\mu F)(144V^2) = 360\mu J$$

Using the first form:

$$U = \frac{1}{2}QV = \frac{1}{2}(60\mu C)(12V) = 360\mu J$$

Using the third form:

$$U = \frac{Q^2}{2C} = \frac{(60\mu C)^2}{2 \times 5\mu F} = \frac{3600\mu C^2}{10\mu F} = 360\mu J$$

All three forms agree: the energy stored is $360\mu J$ (roughly one third of a millijoule).

Making Sense of the Answer: *A capacitor of $5\mu F$ at $12V$ stores about a third of a millijoule, which is in the same ballpark as the kinetic energy of an ant walking at a brisk pace. The stored energy is small in everyday terms, but if it is released through a wire of small resistance in a fraction of a microsecond, the average power dissipated during the discharge can climb into the watt range. Capacitors are not about how much energy is stored; they are about how quickly that energy can be released.*

Think Like a Physicist: *The three forms $U = \frac{1}{2}QV = \frac{1}{2}CV^2 = \frac{Q^2}{2C}$ are the same equation wearing three different costumes, and learning to recognise which costume the question is asking for is half the work of solving energy problems. If the question gives C and V, use the middle form. If the question disconnects the supply and asks what happens next, use the third form,*

because Q is the quantity held fixed and V is the one that changes. Choosing the right form before substituting saves arithmetic and prevents the mistake of using a V that is no longer correct.

BINDER Example 15

A 100 μF capacitor is charged to a potential difference of 200V. The capacitor is then short-circuited through a low-resistance wire and the discharge lasts a total time of 5ms. Find (a) the charge stored on the capacitor before the discharge; (b) the energy stored before the discharge; (c) the average current during the discharge; and (d) the average power delivered to the wire during the discharge.

Solution

(a) The charge stored on the capacitor before the discharge is:

$$Q = CV = (100 \times 10^{-6}\,\text{F})(200\text{V}) = 2 \times 10^{-2}\text{C} = 20\text{mC}$$

(b) The stored energy is:

$$U = \frac{1}{2}CV^2 = \frac{1}{2} \times (100 \times 10^{-6}\,\text{F}) \times (200\text{V})^2 = 2\text{J}$$

(c) The average current during the discharge is the total charge that flows divided by the time:

$$I_{avg} = \frac{Q}{t} = \frac{2 \times 10^{-2}\,\text{C}}{5 \times 10^{-3}\,\text{s}} = 4\text{A}$$

(d) The average power delivered is the total energy divided by the time:

$$P_{avg} = \frac{U}{t} = \frac{2\,\text{J}}{5 \times 10^{-3}\text{s}} = 400\text{W}$$

Making Sense of the Answer: *Two joules sounds modest, but delivered in five milliseconds it amounts to 400 W of average power and four amps of average current, both numbers comparable to those of a small electric kettle, but only for a five-millisecond pulse. The peak current at the start of the discharge is much larger than the four-amp average (because the discharge follows an exponential decay rather than a uniform rate), but the average is the right number to compare with the kettle. A 100 μF capacitor at 200 V is a very ordinary component in a photographic flash unit.*

Think Like a Physicist: *Two of the four quantities asked for here are pure averages (current and power), and average is in both cases the total quantity divided by the time. Average is almost never the same as instantaneous in a capacitor discharge, where current and power both fall exponentially to zero. The student who blurs the two and writes $I_{max} = 4A$ or $P_{max} = 400W$ has misread the question. When a question asks for an average over a discharge, the answer is total energy (or total charge) divided by total time. When a question asks for an initial or peak value, the answer is the current or power at $t = 0$, which is set by the initial voltage and the circuit resistance, not by the time over which the discharge completes.*

REAL Example 16

A capacitor of capacitance $100\mu\text{F}$ is charged to a voltage of 10V by a battery, which is then disconnected. The charged capacitor is now connected by a pair of wires to a second, identical uncharged capacitor of capacitance $100\mu\text{F}$. The two capacitors quickly come to a common voltage. Find (a) the common final voltage; (b) the energy stored on the system before the connection; (c) the energy stored on the system after the connection; (d) account for the difference.

Solution

Mr. Akilikubwa stood at the blackboard with the problem written out, and pointed his chalk in turn at Kipanga and Kipute.

Mr. Akilikubwa: *Four parts, three workings. Kipanga, take part (a). Apply charge conservation and find the final voltage on the combination.*

Kipanga went to the board and wrote out the initial charge on the first capacitor:

$$Q_{\text{initial}} = C_1 V_{1,\text{initial}} = (100\mu\text{F})(10\text{V}) = 1000\mu\text{C}$$

After the connection, the two capacitors are wired in parallel; their equivalent capacitance is:

$$C_{\text{eq}} = C_1 + C_2 = 100\mu\text{F} + 100\mu\text{F} = 200\mu\text{F}$$

Charge cannot disappear, so the same $1000\mu\text{C}$ of total charge is now spread across the parallel combination. The common final voltage is:

$$V_{\text{final}} = \frac{Q_{\text{initial}}}{C_{\text{eq}}} = \frac{1000\mu\text{C}}{200\mu\text{F}} = 5\text{V}$$

Mr. Akilikubwa: *Good. The connection has cut the voltage in half. Kipute, take part (b). The energy before.*

Kipute moved to the board. Before the connection, all the energy is stored on the first capacitor; the second is empty.

$$U_{\text{before}} = \frac{1}{2}C_1 V_{1,\text{initial}}^2 = \frac{1}{2}(100\mu\text{F})(10\text{V})^2 = 5000\mu\text{J} = 5\text{mJ}$$

Mr. Akilikubwa: *And part (c). After the connection.*

Both capacitors now hold the same voltage 5V, so the total stored energy is the sum of the energies of the two:

$$U_{\text{after}} = \frac{1}{2}C_1 V_{\text{final}}^2 + \frac{1}{2}C_2 V_{\text{final}}^2 = \frac{1}{2}(100\mu\text{F})(5\text{V})^2 + \frac{1}{2}(100\mu\text{F})(5\text{V})^2 = 2500\mu\text{J} = 2.5\text{mJ}$$

Kipanga: *Sir, the answer is exactly half the original energy. We started with 5mJ, we now have 2.5mJ. Where is the other half?*

Mr. Akilikubwa: *That is part (d) of the question, and the part the question really cared about. Kipute, you tried this experiment yesterday with the two capacitors on the bench and the LED. What did you observe at the moment of connection?*

Kipute: *There was a small spark at the contact, and a faint click. The wire warmed slightly under my finger.*

Mr. Akilikubwa: *The wire **warmed**, the spark sparked, and a little **sound** was made. That is where the missing 2.5mJ went. The current during the equalisation was enormous (large charge moving in a short interval) and even the small resistance of the wires was enough to dissipate a significant fraction of the stored energy as **heat.** The remainder went into the snap of the spark and the faint **sound wave** you heard as a click.*

Kipanga: *But **what if the wires had been perfect conductors with no resistance at all?** Then no heat could be dissipated, and yet the missing energy is still half of the original.*

Mr. Akilikubwa: *An excellent question, and the answer is that perfect conductors would not save the energy either. A circuit made of two capacitors and a perfect-conducting loop is an oscillator: charge would slosh back and forth between the two capacitors forever, but in doing so it would broadcast **electromagnetic waves** into the surrounding space, as any oscillating charge does. The lost energy is then radiated away as **electromagnetic radiation.** Whether by heat, by sound, by spark, or by radiation, exactly half the original energy must leave the system during a direct capacitor-to-capacitor connection. There is no way to avoid it.*

Kipute: *So if we wanted to transfer energy from one capacitor to another without losing it, a direct wire is not the way.*

Mr. Akilikubwa: *Correct. You would have to put an inductor or a transformer between the two capacitors and let the energy pass through the magnetic field of the inductor on the way. An inductor can store energy without dissipating it (in the ideal case), and the energy can be transferred from one capacitor to the other through the magnetic energy of the inductor with no loss in principle. This is exactly what every camera flash circuit and every defibrillator charger does internally: it puts an inductor in the path between the small battery and the large storage capacitor, precisely to avoid the half-energy loss that a direct connection would suffer.*

Making Sense of the Answer: *The half-energy theorem of capacitor coupling is a feature, not a bug. It explains why every device that needs to take energy from a small supply and pile it into a storage capacitor uses an inductor or a transformer between the two: a direct wire connection would lose half the energy in the very act of transferring it. The numerical result we obtained ($2.5mJ$ out of $5mJ$) is exact, not approximate. It depends only on the fact that the two capacitors were **equal**; for unequal capacitors the lost fraction is also computable and is somewhere between 0 and 0.5.*

Think Like a Physicist: *Whenever an apparent paradox arises in physics, the resolution is almost always found in an assumption that you did not realise you were making. Here, the assumption was that the connecting wires were perfectly idealised: zero resistance, zero inductance, zero radiation. As soon as any one of those three properties is real, and at least one of them always is, the energy bookkeeping closes. The half-energy theorem is a special case of a far more general truth: nature does not let energy disappear, but she is sometimes shy about telling you where it has gone.*

HOT Example 17

A 5μF capacitor is charged to a potential difference of 100V and disconnected from the supply. It is then connected, with like plates joined together, to an uncharged 3μF capacitor. Find (a) the common final voltage on the combination; (b) the charge on each capacitor after the connection; (c) the initial and final stored energies; and (d) the energy lost in making the connection.

Solution

(a) Before the connection, the 5μF capacitor carries:

$$Q_1 = C_1V_1 = (5\mu F)(100V) = 500\mu C$$

The 3μF capacitor carries no charge. The total charge on the joined positive side is therefore 500μC. After the connection, the two capacitors are in **parallel**; their equivalent capacitance is:

$$C_{eq} = C_1 + C_2 = 5\mu F + 3\mu F = 8\mu F$$

The common final voltage is:

$$V_f = \frac{Q_{total}}{C_{eq}} = \frac{500\mu C}{8\mu F} = 62.5V$$

(b) The new charges on the two capacitors at this common voltage:

$$Q_{1,new} = C_1V_f = (5\mu F)(62.5V) = 312.5\mu C$$
$$Q_{2,new} = C_2V_f = (3\mu F)(62.5V) = 187.5\mu C$$

Cross-check on charge conservation: $Q_{1,new} + Q_{2,new} = 312.5\mu C + 187.5\mu C = 500\mu C$, equal to the initial charge, as it must be.

(c) Initial and final stored energies:

$$U_{initial} = \frac{1}{2} C_1 V_i^2 = \frac{1}{2} \times (5 \times 10^{-6}F) \times (100V)^2 = 2.5 \times 10^{-2}J = 25mJ$$

$$U_{final} = \frac{1}{2} C_{eq} V_f^2 = \frac{1}{2} \times (8 \times 10^{-6}F) \times (62.5V)^2 = 1.56 \times 10^{-2}J = 15.6mJ$$

(d) The energy lost in making the connection is the difference:

$$\Delta U_{lost} = U_{initial} - U_{final} = 25mJ - 15.6mJ = 9.4mJ$$

About 9.4 millijoules of the original 25mJ has been dissipated as heat in the connecting wire and as a brief electromagnetic burst.

Making Sense of the Answer: *Roughly 38 per cent of the original stored energy disappears in the very act of connecting two capacitors at different voltages. The loss is large not because the wires were bad, but because any real (or even ideal) connection forces a sudden current through whatever resistance and inductance the loop happens to have, and the work done by that current against itself is exactly the missing energy. This is why energy is never transferred between capacitors through a direct wire in any serious circuit; an inductor or a transformer is always placed in the path so that the energy can pass through a magnetic field and arrive at its destination intact.*

Think Like a Physicist: *Every charge-sharing problem reduces to a single principle: total charge on the joined positive side is unchanged by the connection. Write Q_{total} at the very top of the working, then divide by the **parallel** equivalent capacitance to get V_f. From V_f every other quantity (new charges, new energies, energy lost) follows in one line. Make energy bookkeeping a separate paragraph from the voltage and charge calculation: students who try to compute initial and final energies in the same paragraph almost always lose a factor somewhere.*

REAL Example 18

The photographic flash unit on an old film camera uses a capacitor of capacitance $800\mu F$ charged to a voltage of $300V$. When the camera button is pressed, the capacitor discharges through the flash bulb in $1.0ms$. Calculate (a) the energy stored in the capacitor; (b) the average power dissipated by the flash bulb during the discharge; (c) compare this average power to that of an ordinary $100W$ incandescent lamp.

Solution

(a) The energy stored in the capacitor is:

$$U = \frac{1}{2}CV^2 = \frac{1}{2}(800\mu F)(300V)^2 = \frac{1}{2}(800 \times 10^{-6}F)(90000V^2) = 36J$$

The capacitor stores $36J$ of energy.

(b) The discharge lasts 1.0ms. The average power dissipated by the flash bulb during this brief interval is:

$$P = \frac{U}{t} = \frac{36J}{1.0 \times 10^{-3}s} = 36000W = 36kW$$

(c) Calculating their ratio:

$$\frac{P_{flash}}{P_{bulb}} = \frac{36000W}{100W} = 360$$

The flash bulb operates at 360 times the power of an ordinary $100W$ incandescent lamp for the brief duration of the discharge.

Making Sense of the Answer: *The same energy $36J$ delivered slowly by a $100W$ lamp would take $0.36s$ to be released, an interval over which the human eye sees an ordinary glow. Released instead in $1.0ms$, the same energy produces a flash so intense that the shutter of a camera, which is open for a comparable duration, can record a complete photograph in the time between the start and the end of the discharge. The capacitor has not made any more energy; it has simply concentrated the same energy into a much shorter time, and time-concentration of energy is the entire purpose of the capacitor in a flash circuit.*

Think Like a Physicist: *Whenever a problem asks for the average power dissipated over a discharge interval, the working has two clean steps: first compute the total energy released, and second divide it by the duration. Doing these in the opposite order or trying to compute instantaneous power from current and voltage requires an exponential discharge curve, and is far more work for the same answer. The energy-and-time approach works regardless of the shape of the discharge curve, because U and t are total quantities that do not care about the micro-structure of how the discharge happens. Conservation of energy is a friend on every such problem.*

HOT Example 19

Two capacitors of capacitance $1\mu F$ and $5\mu F$ respectively are charged separately, each to a potential difference of $100V$. The positive plate of the $1\mu F$ capacitor is then wired to the negative plate of the $5\mu F$ capacitor, and the negative plate of the $1\mu F$ capacitor to the positive plate of the $5\mu F$ capacitor. Find (a) the final common voltage on the combination, with its correct sign; (b) the energy lost in the connection.

Solution

(a) Before the connection, each capacitor carries:

$$Q_1 = C_1 V_1 = (1\mu F)(100V) = 100\mu C$$

$$Q_2 = C_2V_2 = (5\mu F)(100V) = 500\mu C$$

Because unlike plates are joined together, the charges of opposite sign partially cancel on the joined side. The net charge on the joined positive side is the magnitude of the difference:

$$Q_{total} = |Q_2 - Q_1| = |500\mu C - 100\mu C| = 400\mu C$$

The sign of the resulting field is determined by the larger of the two initial charges (here Q_2, the charge on the $5\mu F$ capacitor), so the joined plate that was originally the positive plate of the $5\mu F$ capacitor remains positive in the final state.

The two capacitors are now in parallel with $C_{eq} = 1\mu F + 5\mu F = 6\mu F$, and the common final voltage is:

$$V_f = \frac{Q_{total}}{C_{eq}} = \frac{400\mu C}{6\mu F} = 66.7\ V$$

(b) The initial total stored energy is:

$$U_{initial} = \frac{1}{2}C_1V_1^2 + \frac{1}{2}C_2V_2^2 = \frac{1}{2}(1\times10^{-6}F)(100V)^2 + \frac{1}{2}(5\times10^{-6}F)(100V)^2 = 0.03J = 30mJ$$

The final stored energy is:

$$U_{final} = \frac{1}{2}C_{eq}V_f^2 = \frac{1}{2}\times(6\times10^{-6}F)\times(66.7V)^2 = 0.0133J = 13.3\ mJ$$

Energy lost in the connection:

$$\Delta U_{lost} = U_{initial} - U_{final} = 30mJ - 13.3mJ = 16.7mJ$$

Making Sense of the Answer: *More than half of the original stored energy (about 56 per cent) is dissipated during the unlike-plates connection, substantially more than in the like-plates case of the previous example. The reason is geometric. Unlike-plates connection forces a much larger redistribution of charge through the wire than the like-plates connection of the previous example, where charge merely redistributes between two capacitors of like polarity. Larger transient current means more heat dissipated, full stop. The lesson for designers is that unlike-plates accidents are far more damaging than like-plates errors.*

Think Like a Physicist: *Sign conventions are the natural hazard of unlike-plates problems. The safest habit is to compute the magnitudes of all charges first, then ask which initial charge is the larger of the two: that capacitor's polarity determines the sign of the final state, and the magnitude of the difference is the conserved net charge. Trying to carry signed charges through the arithmetic without first clarifying which way they will end up almost always produces sign errors that are difficult to track back to their source.*

HOT Example 20

A parallel-plate capacitor consists of two square metal plates of side $10cm$ separated by an air gap of $1mm$. The capacitor is connected to a $500V$ supply and allowed to charge fully. Calculate the magnitude of the attractive force between the plates.

Solution

The plate area, converted to metres, is:

$$A = (0.10m)^2 = 0.01m^2 = 10^{-2}m^2$$

The gap, converted to metres, is:

$$d = 1mm = 10^{-3}m$$

The field between the plates is $E = V/d$:

$$E = \frac{V}{d} = \frac{500V}{10^{-3}m} = 5\times10^5V/m$$

The magnitude of the attractive force between the plates is $F = \frac{1}{2}\varepsilon_0E^2A$:

$$F = \frac{1}{2}(8.854\times10^{-12}F/m)(5\times10^5V/m)^2(10^{-2}m^2) = 1.1\times10^{-2}N$$

The attractive force between the plates is approximately $1.1\times10^{-2}N$ (about $11mN$).

Making Sense of the Answer: *A force of $11mN$ is small but easily measurable; a sensitive postal balance would detect it without difficulty. The force scales as E^2, so doubling the voltage at the same gap would quadruple the force, and halving the gap at the same voltage would also quadruple it (because $E = V/d$ doubles when d halves). Modern electrostatic loudspeakers, electrostatic micro-electromechanical sensors, and the inkjet printer heads that put writing on paper all exploit this scaling: tiny gaps and modest voltages give very large forces per unit area, in places where mechanical actuators of the same size could not deliver anything useful.*

Think Like a Physicist: *Two formulae for the inter-plate force exist, and they are equivalent: $F = \frac{Q^2}{2A\varepsilon_0}$, which is best when the supply has been disconnected and the charge is held fixed; and $F = \frac{1}{2}\varepsilon_0E^2A$, which is best when the supply is connected and the voltage (and so the field) is held fixed. Picking the right form for the constraint of the problem saves a substitution and often a sign error. As with the energy formula, the discipline of choosing the form before substituting is more than half the battle.*

The energy and the force a capacitor exerts on its own plates are now in hand. The next section turns to a question we have so far avoided: *what difference does it make to all of the above if the gap between the plates is filled with an insulating material instead of air or vacuum?*

DIELECTRICS

So far the gap between the plates has been air, or, if we have wished to be careful about the small ionisation correction, vacuum. In every real capacitor the gap is filled with a solid or liquid insulator: paper, mica, ceramic, polyester, polypropylene, electrolytic gel. The insulator does three jobs at once.

- It increases the capacitance by a factor that depends on the material.
- It raises the voltage at which the capacitor breaks down, often by a large factor.
- It gives the plates something to lean on, so a capacitor the size of a fingernail can have its plates only ten micrometres apart without the smallest mechanical knock putting them in contact.

The insulating filler is called a **dielectric.** *The factor by which the capacitance increases when the dielectric fills the gap* is called the **dielectric constant** of the material, written $\boldsymbol{\varepsilon_r}$. It is dimensionless. Vacuum has $\varepsilon_r = 1$ by definition; air has $\varepsilon_r = 1.0006$, indistinguishable from vacuum in almost every practical calculation. Paper has ε_r between 2 and 4 depending on the grade; mica is around 5 to 7; common ceramic capacitor materials run from 6 up to several thousand; water has $\varepsilon_r \approx 81$; certain titanate ceramics used in modern surface-mount components can reach $\varepsilon_r \approx 10\,000$.

How a dielectric works

Before applying any formula, consider the following figure.

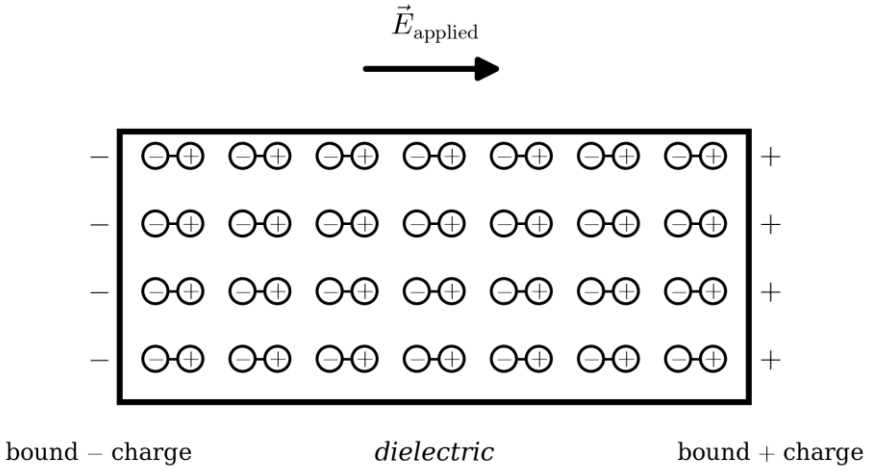


Figure: ***Polar molecules inside a dielectric align with the applied field*** $\boldsymbol{\vec{E}_{applied}}$**.** *The aligned dipoles cancel inside the bulk of the material, leaving uncompensated bound charge only on the two faces of the slab: negative bound charge on the face nearest the positive plate, positive bound charge on the face nearest the negative plate.*

Every molecule contains electrons and nuclei. In the absence of an external field, these charges are arranged so that the molecule has no net dipole moment; or, if it is a polar molecule like water, its many dipoles point in random directions because of thermal motion, and the average dipole is zero. Place the molecule in an electric field, and one of two things happens. If the molecule is non-polar, the electron cloud is pulled slightly against the field direction relative to the positive nuclei, producing a small induced dipole. If the molecule is polar, the existing permanent dipole is partly turned into alignment with the field, against the randomising effect of thermal motion.

Either way, the population of aligned dipoles produces a thin layer of bound positive charge on the face of the dielectric pointing towards the negative plate, and a matching layer of bound negative charge on the face pointing towards the positive plate. These bound charges produce their own electric field, in the opposite direction to the applied field. The net electric field inside the dielectric is therefore **less** than the field that would be there if the dielectric were absent.

Quantitatively, the net electric field inside the dielectric is reduced by exactly the factor ε_r:

$$\mathbf{E_{in\ dielectric}} = \frac{\mathbf{E_{vacuum}}}{\boldsymbol{\varepsilon_r}}$$

For a parallel-plate capacitor with a fixed charge Q on the plates, the field outside the dielectric is set by:

$$\sigma = \frac{Q}{A},$$

and is independent of what fills the gap.

The field inside the dielectric is reduced by ε_r. The voltage across the plates, which is the line integral of the field across the gap, is therefore reduced by ε_r. Since $C = Q/V$, and Q is the same, C is multiplied by ε_r:

$$\mathbf{C = \varepsilon_r C_0}$$

Where C_0 is the capacitance of the same geometry with vacuum (or air) in the gap. The whole of this section reduces, in slogan form, to this one equation. Everything that follows is a matter of unpacking what ε_r means when the dielectric fills only part of the gap, or when more than one dielectric is present.

Two examples follow. The first puts the slogan formula to work; the second uses it to solve a small manufacturing problem.

BINDER Example 21

An air-gap parallel-plate capacitor has capacitance 50pF. A slab of mica with dielectric constant $\varepsilon_r = 6$ is cut to exactly fit the gap and slid into place, filling the space between the plates completely. Find the new capacitance.

Solution

The capacitance with vacuum (or air) in the gap is the value given: $C_0 = 50\,\text{pF}$

With the gap completely filled by a dielectric of constant ε_r, the new capacitance is:

$$C = \varepsilon_r C_0 = 6 \times 50\text{pF} = 300\text{pF}$$

The new capacitance is 300pF.

Making Sense of the Answer: *The mica did exactly what the slogan promised: it multiplied the capacitance by exactly ε_r. Nothing more is needed to predict the new value once ε_r is known. A factor of 6 from a slip of mica thinner than a postage stamp is a remarkable efficiency, and is the reason mica capacitors held the high-precision market for half a century before ceramic took over.*

Think Like a Physicist: *The formula $C = \varepsilon_r C_0$ does not depend on the geometry. It works for any capacitor whose gap is uniformly filled with the same dielectric: parallel plates, concentric cylinders, concentric spheres, anything. The geometric details are hidden inside C_0; the material details are hidden inside ε_r; and the two factors enter the answer multiplicatively.*

REAL Example 22

A small factory in Tabora makes rolled paper capacitors for export. The paper they use is 40mm wide, 0.015mm thick, and has a dielectric constant of 2.5. Two long aluminium foil strips are laid against the paper, a second sheet of paper goes on top, and the four-layer sandwich is rolled into a small cylinder. The capacitor is meant to have a capacitance of 2.0μF.

(a) Find the length of paper, and so of foil, that the factory must roll up to make a single capacitor of this rating.
(b) Explain one practical reason why the rolled construction works, and one practical reason why the design might fail in service.

Solution

Treat the unrolled sandwich as a single parallel-plate capacitor of plate area $A = Lw$, where L is the length of the foil and $w = 40\,\text{mm} = 0.040\,\text{m}$ is the width, with the paper as the dielectric of thickness $d = 0.015\text{mm} = 1.5 \times 10^{-5}\,\text{m}$. The capacitance is:

$$C = \frac{\varepsilon_r \varepsilon_0 A}{d} = \frac{\varepsilon_r \varepsilon_0 Lw}{d}$$

Solving for L:

$$L = \frac{Cd}{\varepsilon_r \varepsilon_0 w} = \frac{(2.0 \times 10^{-6}\,\text{F})(1.5 \times 10^{-5}\,\text{m})}{(2.5)(8.854 \times 10^{-12}\,\text{F/m})(0.040\,\text{m})} = 33.9\,\text{m}$$

A length of about 33.9m of paper, and the same of each foil strip, is required for one 2.0μF capacitor.

(b) Practical observations.

Why the construction works: The rolled-strip design delivers a very large plate area in a very small package. Thirty-four metres of foil is the area that the formula demands for a 2μF capacitor at this gap and dielectric, but a 34m-long flat sheet would not fit on a workbench. Rolling the sandwich into a cylinder a few centimetres long preserves the foil-to-foil area exactly (the formula does not care how the area is shaped) while shrinking the package down to something that fits comfortably onto a circuit board. The rolled construction is therefore a geometric trick that buys area density without changing the electrical parameters.

Why the design might fail in service: Paper is hygroscopic; it absorbs water vapour from the surrounding air. Once moisture has soaked into the paper between the foils, the dielectric strength of the paper drops sharply, because dissolved ions in the absorbed water provide a conduction path between the plates. In a humid coastal climate such as that of Tanga or Dar es Salaam, an unsealed paper capacitor can lose its rated voltage in months. Manufacturers solve this by impregnating the paper with oil or wax and sealing the rolled cylinder inside a metal can; an exported capacitor that arrives without such sealing will eventually short out in service.

Making Sense of the Answer: *Thirty-four metres of paper to make a single small component sounds extravagant until one remembers that the rolled cylinder is perhaps three centimetres long. The roll squeezes that long strip into a package the size*

of a thumb. The trick of the rolled construction is exactly this: trading a long thin sheet for a stubby cylinder, with no change in the area or in the gap that the formula sees.

Think Like a Physicist: *Notice how nothing in the calculation depended on the radius of the rolled cylinder, on how tightly the cylinder was wound, or on how many turns the paper made. All that matters is the total area of foil facing foil with paper between them, and that is* Lw, *the area of the unrolled sandwich. When a problem looks geometrically intimidating, ask yourself whether the answer cares about the intimidating part of the geometry. Very often, it does not.*

Partially-filled gap

The mica slab in the last example was cut to exactly fit the gap. Suppose now it is too thin: it fills part of the gap but leaves a small air space between the slab and one of the plates. Or suppose, for safety reasons, the designer wants a sliver of air on either side of the slab so that the slab is not in electrical contact with either plate. *What is the capacitance of such a partially-filled gap?*

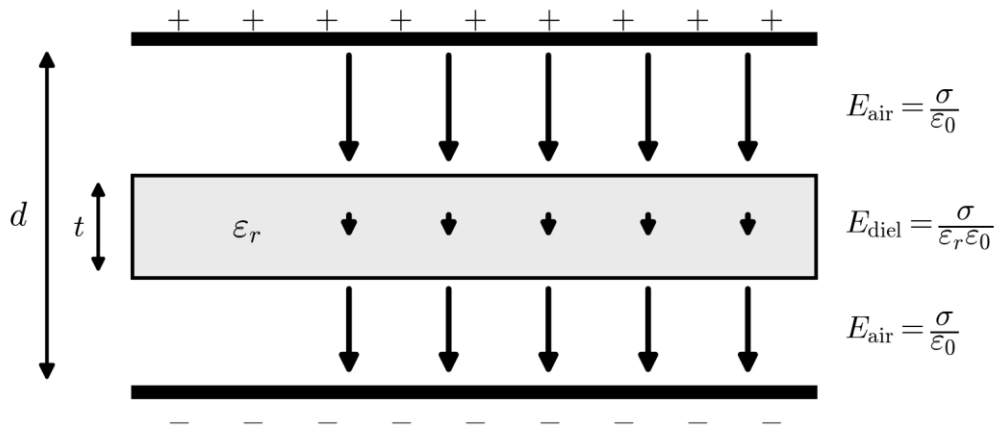


Figure: ***Partially-filled gap.*** *Total separation d, slab of thickness t with dielectric constant ε_r somewhere in the gap, leaving $(d - t)$ of air. The field in the air region is the full vacuum value σ/ε_0; the field inside the dielectric is reduced to $\sigma/(\varepsilon_r\varepsilon_0)$. The arrows inside the slab are drawn shorter to indicate the smaller field strength there.*

Let the total gap be d and the slab thickness be t, with $t < d$. The plates carry surface charge density $\sigma = Q/A$. The field outside the dielectric is the full vacuum value, σ/ε_0; the field inside the dielectric is reduced by ε_r to $\sigma/(\varepsilon_r\varepsilon_0)$. The potential difference across the whole gap is the sum of the two contributions:

$$V = E_{\text{air}}(d - t) + E_{\text{diel}}\, t = \frac{\sigma}{\varepsilon_0}(d - t) + \frac{\sigma}{\varepsilon_r \varepsilon_0}\, t$$

Factor out σ/ε_0:

$$V = \frac{\sigma}{\varepsilon_0}\left((d - t) + \frac{t}{\varepsilon_r}\right) = \frac{Q}{\varepsilon_0 A}\left((d - t) + \frac{t}{\varepsilon_r}\right)$$

The capacitance is Q/V:

$$\boldsymbol{C = \frac{\varepsilon_0 A}{(d - t) + \frac{t}{\varepsilon_r}}}$$

Two limit checks.

- When $t \to 0$, the slab vanishes and the bracket reduces to d, giving $C = \frac{\varepsilon_0 A}{d}$, the full air-gap capacitance, as it must.
- When $t \to d$, the slab fills the gap and the bracket reduces to d/ε_r, giving $C = \frac{\varepsilon_r \varepsilon_0 A}{d}$, the full dielectric result. Both limits recover the cases we already understand.

A second way to arrive at the same formula is to notice that the partially-filled gap is two capacitors in series: an air capacitor of thickness $(d - t)$ stacked on top of a dielectric capacitor of thickness t. The series combination gives the same answer, with a little less arithmetic if the student is comfortable with series-capacitor algebra. Either route is acceptable.

Two examples follow.

HOT Example 23

A parallel-plate capacitor has plates of area 200cm² separated by an air gap of 0.5mm. A slab of dielectric constant 4 and thickness 0.3mm is inserted, with the remaining 0.2mm left as air. Find the new capacitance.

Solution

Convert all lengths and the area to SI base units:

$$A = 200\text{cm}^2 = 2.00 \times 10^{-2}\text{m}^2 \quad d = 0.5\text{mm} = 5 \times 10^{-4}\text{m} \quad t = 0.3\text{mm} = 3 \times 10^{-4}\text{m} \quad \varepsilon_r = 4$$

Use the partial-fill formula:

$$C = \frac{\varepsilon_0 A}{(d-t)+\frac{t}{\varepsilon_r}} = \frac{(8.854\times10^{-12}\,\text{F/m})(2.00\times10^{-2}\,\text{m}^2)}{(5\times10^{-4}\,\text{m}-3\times10^{-4}\,\text{m})+\left(\frac{3\times10^{-4}\,\text{m}}{4}\right)} = 6.44\times10^{-10}\,\text{F} = 644\,\text{pF}$$

The partially-filled capacitor stores about 644pF.

Alternative solution (as two capacitors in series)

The partial-fill geometry is two capacitors stacked between the same pair of plates: an air-only layer of thickness $(d-t) = 0.2\text{mm}$ on one side of the slab, and a fully dielectric-filled layer of thickness $t = 0.3\text{mm}$ on the other. The same charge Q passes through both layers, so they are in series.

The air layer alone has capacitance:

$$C_{air} = \frac{\varepsilon_0 A}{d-t} = \frac{(8.854\times10^{-12}\,\text{F/m})(2.00\times10^{-2}\,\text{m}^2)}{2\times10^{-4}\,\text{m}} = 8.854\times10^{-10}\,\text{F} = 885.4\text{pF}$$

The dielectric layer alone has capacitance:

$$C_{diel} = \frac{\varepsilon_r\varepsilon_0 A}{t} = \frac{(4)(8.854\times10^{-12}\text{F/m})(2.00\times10^{-2}\text{m}^2)}{3\times10^{-4}\text{m}} = 2.361\times10^{-9}\text{F} = 2361\text{pF}$$

The product-over-sum form of the two-capacitor series rule gives:

$$C = \frac{C_{air}\,C_{diel}}{C_{air}+C_{diel}} = \frac{(885.4\,\text{pF})(2361\,\text{pF})}{885.4\,\text{pF}+2361\,\text{pF}} = 644\,\text{pF}$$

Either route is acceptable in an examination.

Making Sense of the Answer: *Compare with the two endpoints. The same gap filled entirely with air would give $C_0 = \varepsilon_0 A/d \approx 354\,pF$; filled entirely with the $\varepsilon_r = 4$ slab it would give $4C_0 \approx 1416\,pF$. The actual partial-fill answer of 644pF sits between these two, closer to the air endpoint than to the full-slab endpoint. This is the correct intuition: 0.3mm of slab in a 0.5mm gap is more than half the gap, but the air strip is still close enough to the full-air answer that the boost from the slab is only about 82 per cent of the maximum.*

Think Like a Physicist: *A slab in a gap is two capacitors in series, an air one and a dielectric one. Series capacitors are dominated by the smaller of the two: whichever capacitor has the smaller C controls the total. The air strip here has $C_{air} = \varepsilon_0 A/(d-t) \approx 885\,pF$, and the slab itself has $C_{slab} = \varepsilon_r\varepsilon_0 A/t \approx 2360\,pF$. The air strip is the smaller and therefore dominates. To raise C further, the designer would have to thin the air strip, which is exactly why real commercial capacitors leave no air at all and fill the entire gap with dielectric.*

HOT Example 24

A parallel-plate capacitor is charged from a battery, and the battery is then disconnected. A slab of dielectric of thickness 3mm is then introduced into the gap. To keep the potential difference across the plates the same as it was before the slab was inserted, the plate separation must be increased by 2.4mm. Find the dielectric constant of the slab.

Solution

After the battery is disconnected, the charge Q on the plates is fixed. The potential difference $V = Q/C$ therefore stays the same if and only if the capacitance C stays the same. The condition of the problem is therefore $C_{new} = C_{old}$.

Before the slab is inserted, the capacitor is a pure air gap of width d, with capacitance:

$$C_{old} = \frac{\varepsilon_0 A}{d}$$

After the slab is inserted and the gap widened by $\Delta d = 2.4\,\text{mm}$, the new total gap is $(d+\Delta d)$, of which thickness $t = 3\,\text{mm}$ is slab and the rest is air. Apply the partial-fill formula:

$$C_{new} = \frac{\varepsilon_0 A}{(d+\Delta d-t)+\frac{t}{\varepsilon_r}}$$

Setting $C_{new} = C_{old}$, the $\varepsilon_0 A$ factors match and the numerators equate, so the denominators must equate:

$$d = (d+\Delta d-t)+\frac{t}{\varepsilon_r}$$

The d terms cancel:

$$0 = \Delta d - t + \frac{t}{\varepsilon_r}$$

Solve for ε_r:

$$\varepsilon_r = \frac{t}{t-\Delta d} = \frac{3\text{mm}}{3\text{mm}-2.4\text{mm}} = 5$$

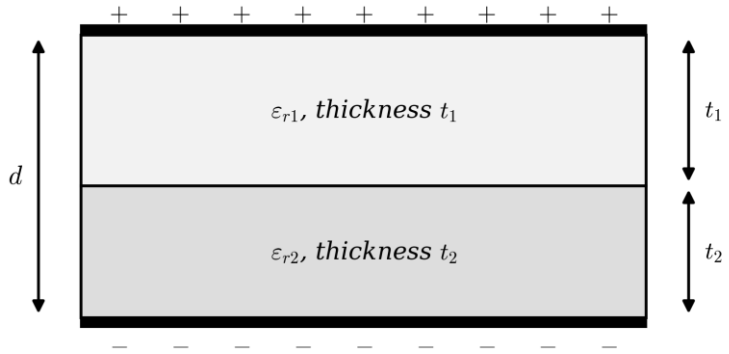
The dielectric constant of the slab is 5.

Making Sense of the Answer: $\varepsilon_r = 5$ *is in the ballpark of mica, glass, or a stiff polymer; not the highest-K materials, but well above paper. The arithmetic is striking in how little of it survives. The answer comes from a single relation,* $t(1 - 1/\varepsilon_r) = \Delta d$*, which has the simple interpretation that the slab replaces a layer of air of thickness* Δd*, and the dielectric constant is set by how thick a slab is required to replace a given thickness of air.*

Think Like a Physicist: *Two-line solutions are not always easier than three-line ones, but they always teach more. The clean relation* $t(1 - 1/\varepsilon_r) = \Delta d$ *works for any slab and any compensation distance, and is worth keeping. It says that a thick slab with high* ε_r *is hardly any electrically denser than a thin slab with the same effect, because the ratio* $\Delta d/t$ *saturates at 1 as* $\varepsilon_r \to \infty$*. A perfect conductor in the gap (where* $\varepsilon_r \to \infty$*) replaces its own thickness exactly, which is the result we will meet again later in this chapter for the conducting slab.*

Two dielectrics in one gap

If a single slab fills the gap, the answer is $C = \varepsilon_r \varepsilon_0 A/d$. If the gap is only partly filled, the answer is the formula derived just now. But sometimes a designer wants two different materials between the same pair of plates, either stacked one on top of the other across the full plate area, or sitting side by side, each covering its own piece of the plate area. Each arrangement gives a different formula, and which formula applies depends entirely on whether the two slabs share the same charge or the same voltage.



stacked configuration: same Q, voltages add (series)

Figure: ***Stacked configuration.*** *Two slabs of dielectric constants* ε_{r1} *and* ε_{r2}*, of thicknesses* t_1 *and* t_2*, both spanning the full plate area A, stacked to fill the gap* $d = t_1 + t_2$*. The same charge Q passes through both slabs, so they are in series.*

In the stacked configuration, both slabs span the full area A but each occupies a different fraction of the gap. The same charge Q flows through both regions, since charge cannot accumulate at the slab-to-slab interface; whatever leaves one plate must arrive at the other, passing through both materials in succession. The voltage across the pair is the sum of the voltages across each slab. This is the algebra of capacitors in series:

$$\frac{1}{C} = \frac{1}{C_1} + \frac{1}{C_2} = \frac{t_1}{\varepsilon_{r1}\varepsilon_0 A} + \frac{t_2}{\varepsilon_{r2}\varepsilon_0 A} = \frac{\varepsilon_{r2}t_1 + \varepsilon_{r1}t_2}{\varepsilon_{r1}\varepsilon_{r2}\varepsilon_0 A}$$

$$\mathbf{C = \frac{\varepsilon_{r1}\varepsilon_{r2}\varepsilon_0 A}{\varepsilon_{r2}t_1 + \varepsilon_{r1}t_2} = \frac{\varepsilon_0 A}{\frac{t_1}{\varepsilon_{r1}} + \frac{t_2}{\varepsilon_{r2}}}}$$

For three or more layers the formula extends in the obvious way: *the denominator gains another* t/ε_r *term for each new layer.* In the symmetric case of two equal layers each of thickness d/2:

$$\mathbf{C_{stacked,\,equal} = \frac{2\varepsilon_0 A}{d}\left(\frac{\varepsilon_{r1}\varepsilon_{r2}}{\varepsilon_{r1} + \varepsilon_{r2}}\right)}$$

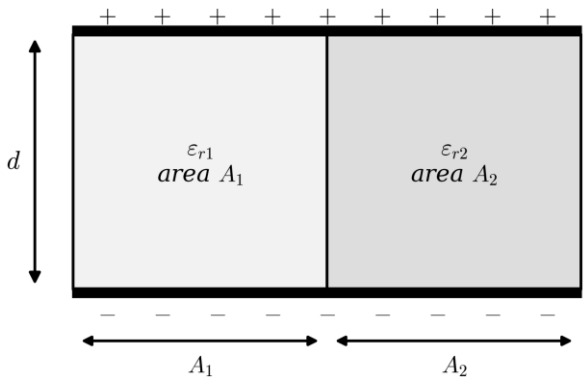


Figure: ***Side-by-side configuration.*** *Two slabs of dielectric constants ε_{r1} and ε_{r2}, each spanning the full gap d but each covering a different piece of the plate area, A_1 and A_2, with $A_1 + A_2 = A$. Both halves see the same voltage V, so they are in parallel.*

In the side-by-side configuration, both slabs span the full gap d but each covers a different fraction of the plate area. Both halves see the same voltage V across the gap; the total charge is the sum of the charges on each half. This is the algebra of capacitors in parallel:

$$\mathbf{C} = C_1 + C_2 = \frac{\varepsilon_{r1}\varepsilon_0 A_1}{d} + \frac{\varepsilon_{r2}\varepsilon_0 A_2}{d} = \frac{\boldsymbol{\varepsilon_0}}{\mathbf{d}}(\boldsymbol{\varepsilon_{r1}A_1 + \varepsilon_{r2}A_2})$$

For two slabs covering equal halves of the plate, $A_1 = A_2 = \frac{A}{2}$:

$$\mathbf{C_{side\text{-}by\text{-}side,\,equal}} = \frac{\boldsymbol{\varepsilon_0 A}}{\mathbf{2\,d}}(\boldsymbol{\varepsilon_{r1} + \varepsilon_{r2}})$$

The two configurations give different formulae for the same pair of materials. The stacked formula uses the harmonic mean of the dielectric constants, weighted by thickness; the side-by-side formula uses the arithmetic mean, weighted by area. The arithmetic mean is always larger than the harmonic mean, so for the same materials the side-by-side configuration always gives the larger capacitance, a result worth checking by direct substitution after every calculation.

These two configurations are nothing more than the series and parallel combinations from earlier in the chapter, dressed in different clothes. The dielectric slabs are themselves the capacitors. Once the student has spotted which configuration applies, the rest is bookkeeping. One example to demonstrate.

HOT Example 25

A parallel-plate capacitor has plates of area 100cm² separated by a gap of 4mm. The gap is completely filled by two dielectric slabs of equal thickness 2mm each. The upper slab has dielectric constant $\varepsilon_{r1} = 2$ and the lower slab has $\varepsilon_{r2} = 6$. Each slab spans the full plate area. Find the capacitance.

Solution

The two slabs are stacked one above the other, each spanning the full plate area. This is the stacked, or series, configuration. Use the symmetric form of the stacked formula, since $t_1 = t_2 = d/2$:

$$C = \frac{2\varepsilon_0 A}{d}\left(\frac{\varepsilon_{r1}\varepsilon_{r2}}{\varepsilon_{r1} + \varepsilon_{r2}}\right)$$

$$C = \frac{2(8.854 \times 10^{-12}\ \text{F/m})(1.00 \times 10^{-2}\ \text{m}^2)}{4 \times 10^{-3}\ \text{m}}\left(\frac{2 \times 6}{2 + 6}\right) = 66.4\ \text{pF}$$

The stacked two-dielectric capacitor stores about 66.4pF.

Making Sense of the Answer: *Compare with the two materials acting alone in the full gap. A full $\varepsilon_r = 2$ fill would give $2C_0 = 44.3\,pF$; a full $\varepsilon_r = 6$ fill would give $6C_0 = 132.8\,pF$. The stacked answer of 66.4pF sits between these but well below the arithmetic mean of 88.5pF, because series combinations are dominated by the smaller capacitor, which here is the $\varepsilon_r = 2$ slab. If the same two slabs were arranged side by side instead, the answer would be $\varepsilon_0 A(2 + 6)/(2d) = 88.5\,pF$, exactly the arithmetic-mean prediction and larger than the stacked result.*

Think Like a Physicist: *Faced with a two-dielectric problem, the very first question to ask is not "what formula?" but "what is conserved across the slab boundary?". If the same charge has to pass through both slabs (stacked), the formula is series. If the same voltage sits across both slabs (side by side), the formula is parallel. The geometry tells you which conservation applies, and the conservation tells you the formula. The harmonic-versus-arithmetic mean distinction follows automatically.*

The capacitor and its materials are now in hand. The next part of this section asks what happens when the dielectric is inserted into a **charged** capacitor, with the battery either connected or disconnected, and the surprising contrast between those two cases.

Battery connected versus battery disconnected

When a dielectric is inserted into a charged capacitor, the result depends decisively on whether the battery is still connected to the plates or has been removed. The two cases give different final voltages, different final stored charges, different final field strengths, and different final stored energies. Examination problems frequently put this contrast at the centre of a question, and students who try to answer them without first sorting out which case applies almost always get the signs wrong. The contrast is worth pinning down once, carefully, and then using as a checklist.

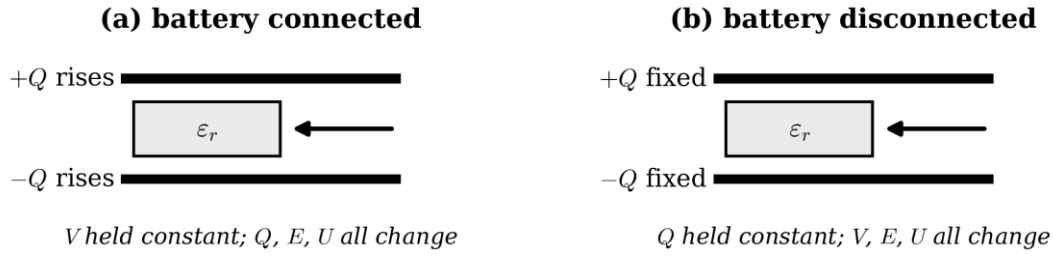


Figure: ***A dielectric slab is being inserted into a parallel-plate capacitor.*** *In (a) the battery is connected throughout the insertion, so the voltage across the plates is fixed by the battery and the charge on the plates is whatever the battery has to supply to satisfy the new capacitance. In (b) the battery has been disconnected before the slab is introduced, so the charge on the plates is locked in and the voltage adjusts to whatever the new capacitance requires.*

Take both cases starting from the same initial state: an air-gap capacitor of capacitance C_0 charged to a voltage V_0, holding a charge $Q_0 = C_0V_0$, with a field $E_0 = V_0/d$ between the plates and a stored energy $U_0 = 1/2\, C_0V_0^2$. A slab of dielectric constant ε_r is now inserted to fill the gap completely.

Case 1: battery connected

The battery enforces $V = V_0$ at all times. The capacitance rises to ε_rC_0, so the charge on the plates must rise to $Q = (\varepsilon_rC_0)V_0 = \varepsilon_rQ_0$.
The field: Since V is unchanged and d is unchanged, the electric field stays at E_0.

The stored energy: $1/2\, CV^2$, rises to $1/2\,(\varepsilon_rC_0)V_0^2 = \varepsilon_rU_0$. The extra charge $\varepsilon_rQ_0 - Q_0 = (\varepsilon_r - 1)Q_0$ is supplied by the battery, which does work $(\varepsilon_r - 1)Q_0V_0$ in the process. Of this work, the capacitor stores an extra $(\varepsilon_r - 1)U_0$; the rest, also $(\varepsilon_r - 1)U_0$, is done as mechanical work on the dielectric drawing it into the gap, against the fringing-field forces at the edges of the plates. The energy bookkeeping balances cleanly.

Case 2: battery disconnected

The charge $Q = Q_0$ is locked in. The capacitance still rises to ε_rC_0, but with Q fixed and C larger, the voltage must drop. From $V = \frac{Q}{C}$, the new voltage is $V = \frac{Q_0}{\varepsilon_rC_0} = \frac{V_0}{\varepsilon_r}$.
The field: Since the gap is unchanged but V has dropped, the electric field falls to $E = \frac{V}{d} = \frac{E_0}{\varepsilon_r}$.
The stored energy: $\frac{Q^2}{2C}$, drops to $\frac{Q_0^2}{2\varepsilon_rC_0} = \frac{U_0}{\varepsilon_r}$. The energy lost, $U_0\left(1 - \frac{1}{\varepsilon_r}\right)$, is also done as mechanical work on the dielectric drawing it into the gap; it comes out of the capacitor's stored energy rather than out of a battery, because no battery is present.

Both cases lose energy from the field-and-battery system. Both lose it for the same physical reason: *the dielectric is drawn into the gap and the system does work on it.* The difference between the two cases is purely accounting. In case 1, the battery refills the energy account; in case 2, the account starts smaller because nothing refills it.

The five-row summary, with every quantity tracked side by side, is the table every student should keep in mind while solving a problem of this type.

Quantity	Initial (air)	Battery connected	Battery disconnected
Capacitance	C_0	ε_rC_0	ε_rC_0
Charge	Q_0	ε_rQ_0	Q_0
Voltage	V_0	V_0	$\frac{V_0}{\varepsilon_r}$
Electric field	E_0	E_0	$\frac{E_0}{\varepsilon_r}$
Stored energy	U_0	ε_rU_0	$\frac{U_0}{\varepsilon_r}$

One worked example follows, with numbers.

HOT Example 26

An air-gap parallel-plate capacitor has capacitance 8pF, and is charged from a battery to a potential difference of 200V. The dielectric in the gap is then replaced by a mica slab of dielectric constant $\varepsilon_r = 4$. Compute the capacitance, charge, voltage, and stored energy under two scenarios: (a) the battery is connected throughout the replacement; (b) the battery has been disconnected before the replacement.

Solution

Compute the initial quantities first; both cases share them.

The initial capacitance is the value given:

$$C_0 = 8\text{pF} = 8 \times 10^{-12}\text{F}$$

The initial charge:

$$Q_0 = C_0 V_0 = (8 \times 10^{-12}\text{F})(200\text{V}) = 1.6 \times 10^{-9}\text{C} = 1.6\text{nC}$$

The initial stored energy:

$$U_0 = 1/2\, C_0 V_0^2 = 1/2\, (8 \times 10^{-12}\text{F})(200\text{V})^2 = 1.6 \times 10^{-7}\text{J} = 0.16\mu\text{J}$$

The new capacitance is the same in both cases:

$$C = \varepsilon_r C_0 = (4)(8\text{pF}) = 32\text{pF}$$

(a) battery connected:

V is held at $V_0 = 200\text{V}$. The new charge:

$$Q = CV_0 = (32 \times 10^{-12}\text{F})(200\text{V}) = 6.4 \times 10^{-9}\text{C} = 6.4\text{nC}$$

The new stored energy:

$$U = 1/2\, CV_0^2 = 1/2\, (32 \times 10^{-12}\text{F})(200\text{V})^2 = 6.4 \times 10^{-7}\text{J} = 0.64\mu\text{J}$$

The charge has risen by a factor of $\varepsilon_r = 4$; the energy has risen by a factor of $\varepsilon_r = 4$.

(b) battery disconnected: Q is held at $Q_0 = 1.6\text{nC}$. The new voltage:

$$V = \frac{Q_0}{C} = \frac{1.6 \times 10^{-9}\text{C}}{32 \times 10^{-12}\text{F}} = 50\text{V}$$

The new stored energy:

$$U = \frac{Q_0^2}{2C} = \frac{(1.6 \times 10^{-9}\text{C})^2}{2(32 \times 10^{-12}\text{F})} = 4 \times 10^{-8}\text{J} = 0.04\mu\text{J}$$

The voltage has dropped by a factor of $\varepsilon_r = 4$; the energy has dropped by a factor of $\varepsilon_r = 4$.

Both columns of the comparison are now filled with numbers.

(a) ends at 6.4nC, 200V, 0.64μJ;

(b) ends at 1.6nC, 50V, 0.04μJ.

The two final energies (0.64μJ and 0.04μJ) have geometric mean $\sqrt{(0.64)(0.04)}\mu\text{J} = 0.16\mu\text{J}$, which is the initial energy, as expected.

Making Sense of the Answer: *The two cases differ by a factor of sixteen in the stored energy. In (a) the capacitor ends up holding sixteen times the energy that (b) ends up holding. The reason is that (a) had a battery doing extra work for the duration of the insertion, while (b) had no such top-up; the battery in (a) supplied four times as much energy as case (b) had at the start, and a quarter of that supply went into stored field while the rest went into mechanical work on the slab.*

Think Like a Physicist: *Always begin a battery-connected-vs-disconnected problem by writing down the five-row table for the geometry, even if some of the rows are not asked. Filling the table forces the student to commit to a single column and stops the typical error of mixing rows from both columns. The geometric-mean check on the two final energies costs almost no time and catches arithmetic slips in real time. Examination question that asks for "the energy lost" almost always expect the student to compute the difference between the initial energy and the final energy of the appropriate column, not the difference between the two columns.*

Dielectric constant, dielectric strength, water, and temperature

So far the dielectric has been treated as if it were specified by one number, the dielectric constant ε_r. Two further properties matter in practice. The first is the maximum field the material can withstand before it breaks down and conducts; this is called the **dielectric strength** of the material, measured in volts per metre. The second is how the dielectric constant changes with temperature; for most polar materials, ε_r drops as the temperature rises, sometimes by a substantial fraction.

Dielectric strength is a different quantity from dielectric constant, and the two are not correlated. A high-ε_r material may have a low dielectric strength, in which case it cannot hold a large voltage even though it would, in principle, store a lot of charge for the voltage it can hold. A low-ε_r material with a high dielectric strength may end up storing more energy than the high-ε_r material, because the maximum field, and so the maximum energy density, depends on the square of the field. The following short table gives representative values to fix the orders of magnitude.

Material	**Dielectric constant ε_r**	**Dielectric strength**
Vacuum	1 (exact)	∞
Dry air	1.0006	3MV/m
Paper	2 to 4	16MV/m
Polyethylene	2.3	50MV/m
Polystyrene	2.6	25MV/m
Mica	5 to 7	100MV/m
Glass	5 to 10	14MV/m
Water (deionised)	80	low (depends on ions)
Titanate ceramic	up to 10000	varies, often modest

Water is a striking entry in this list. Its dielectric constant is enormous, around 80 at room temperature, far higher than any common solid. It would seem to be the perfect dielectric. In practice, water is almost never used as a capacitor dielectric, for three reasons.

- Water dissolves ions readily; even highly purified water has enough dissolved ionic content to conduct measurably, which short-circuits any capacitor.
- The dielectric strength of water depends sharply on the ionic content, and trace contamination drops it by orders of magnitude.
- The dielectric constant of water is strongly temperature-dependent: ε_r drops from about 88 at $0°\text{C}$ to about 56 at $100°\text{C}$, a fall of more than a third over the range of liquid water. A capacitor whose capacitance changed by a third over the operating temperature would be useless for any precision application.

The temperature dependence of polar liquids has a simple physical reason. At higher temperatures, the random thermal motion of the molecules competes with the aligning effect of the applied field. The hotter the liquid, the more thoroughly thermal motion scrambles the dipole orientation, and the smaller the net polarisation produced by the field. Since the polarisation is what gives a material its ε_r above 1, a smaller polarisation means a smaller ε_r. Non-polar dielectrics (paper, polyethylene, mica) show much weaker temperature dependence because their ε_r comes from induced dipoles rather than the alignment of permanent ones, and induced dipoles are largely insensitive to temperature.

Two examples follow. The first is a quick check on the figure-of-merit for energy storage; the second is a real-world drift problem with a soil moisture sensor in Morogoro.

BINDER Example 27

A capacitor designer must choose between two candidate dielectric materials for a high-energy storage capacitor. Material A has dielectric constant $\varepsilon_r = 100$ and dielectric strength 5MV/m. Material B has $\varepsilon_r = 5$ and dielectric strength 200MV/m. For a fixed gap and fixed plate area, which material stores more energy at its breakdown limit, and by what factor?

Solution

The maximum energy density in a dielectric capacitor is the energy density of the field, evaluated at the breakdown field:

$$u_{\max} = 1/2\,\varepsilon_r \varepsilon_0 E_{\text{strength}}^2$$

For Material A:

$$u_A = 1/2\,(100)(8.854 \times 10^{-12}\text{F/m})(5 \times 10^6\text{V/m})^2 = 1.107 \times 10^4\text{J/m}^3$$

For Material B:

$$u_B = 1/2\,(5)(8.854 \times 10^{-12}\text{F/m})(2 \times 10^8\text{V/m})^2 = 8.854 \times 10^5\text{J/m}^3$$

The ratio of maximum stored energies, for the same gap and area, is:

$$\frac{u_B}{u_A} = \frac{8.854 \times 10^5\text{J/m}^3}{1.107 \times 10^4\text{J/m}^3} = 80$$

Material B stores 80 times more energy at the breakdown limit than Material A does, despite Material A having a dielectric constant twenty times larger.

Making Sense of the Answer: *The figure of merit for energy storage is not ε_r alone but the product $\varepsilon_r E_{strength}^2$. Material B's strength is 40 times larger than Material A's, so its E^2 contribution is $40^2 = 1600$ times larger; Material A's ε_r is 20 times larger, which means B is down by a factor of 20 in the ε_r contribution. The net advantage for B is $1600/20 = 80$, and that is exactly the factor by which B beats A in energy density. This is why modern high-energy-density capacitors are built around*

polymer films with modest ε_r but very high dielectric strength, rather than around the high-ε_r titanate ceramics that dominate signal-coupling applications.

Think Like a Physicist: *Whenever a quantity is bounded by something that enters squared in the figure of merit, the bound dominates. Maximum voltage on a capacitor, maximum stress on a beam, maximum current in a resistor: in every case the limit on one factor that appears squared in the answer matters more than a much larger first-power factor that does not face such a limit. The lesson generalises: hunt for the squared variable first when comparing designs.*

REAL Example 28

In a soil-moisture monitoring station at Sokoine University of Agriculture (SUA) in Morogoro, a water-dielectric capacitor sensor is calibrated at $25°C$, where the dielectric constant of the water in the sensor is 78. During the dry-season midday peak the sensor temperature rises to $45°C$, where the dielectric constant of the water has dropped to 71. The sensor electronics convert the measured capacitance directly into a soil moisture reading, with no temperature compensation. Find the percentage error in the moisture reading caused purely by the temperature drift, and state the sign of the error.

Solution

The capacitance of the sensor is proportional to the dielectric constant of the water in it, with all geometric factors held fixed. The fractional change in capacitance, from the calibration temperature to the operating temperature, is:

$$\frac{\Delta C}{C_{25}} = \frac{C_{45} - C_{25}}{C_{25}} = \frac{\varepsilon_{r,45} - \varepsilon_{r,25}}{\varepsilon_{r,25}} = \frac{71 - 78}{78} = \frac{-7}{78} = -0.0897 \text{ (negative)}$$

The percentage error is approximately:

$$\left|\frac{\Delta C}{C_{25}}\right| \times 100\% = 8.97\% \quad \text{(reading low)}$$

The capacitance has fallen by about 9 per cent. Since the electronics read off moisture content as proportional to capacitance, the moisture reading falls by the same 9 per cent. The error is **negative**: the sensor reports the soil as drier than it actually is.

Making Sense of the Answer: *A nine per cent error is small enough to miss in a single reading but large enough to skew a season-long irrigation schedule. The error has a definite sign: it always tells the operator the soil is drier than it really is, on a hot afternoon. An operator who responds to such readings by ordering extra irrigation will over-water on hot days and under-water on cool days, a pattern that disrupts both the soil moisture and the irrigation budget. The fix is either a temperature-compensated electronics design, in which the reading is corrected using a separate temperature sensor, or a non-polar dielectric like a polymer film whose ε_r changes much less with temperature.*

Think Like a Physicist: *Whenever a transducer relies on a material property that has a temperature dependence, the temperature dependence enters the calibration. A device that reads correctly in the laboratory at one temperature need not read correctly in the field at a different temperature. The systematic, signed nature of the error (always reading low when it gets hotter) is a strong clue to its origin: random errors do not have a definite sign; systematic errors do.*

Capacitor partly immersed in a dielectric liquid

Before introducing any formula, let us first examine an important physical situation involving a dielectric liquid and a charged capacitor.

Imagine a vertical parallel-plate capacitor lowered partially into a tank containing a dielectric liquid. The lower portion of the plates lies below the liquid surface, while the upper portion remains above it. The plates are connected to a voltage source, so a constant potential difference is maintained across them.

One might expect the liquid level between the plates to remain the same as the level in the surrounding tank. Surprisingly, this does not happen. Instead, the liquid climbs upward between the plates and settles at a height h above the external liquid surface.

This rise is not caused by capillary action, but by electrostatic forces. The electric field between the plates polarises the dielectric liquid, and the fringing electric field near the boundary pulls the liquid into the region where the electric field exists. In effect, the capacitor attracts the dielectric into the gap in much the same way that it would pull in a solid dielectric slab. The liquid rises until the upward electric force is exactly balanced by the downward gravitational force of the raised liquid column, establishing equilibrium.

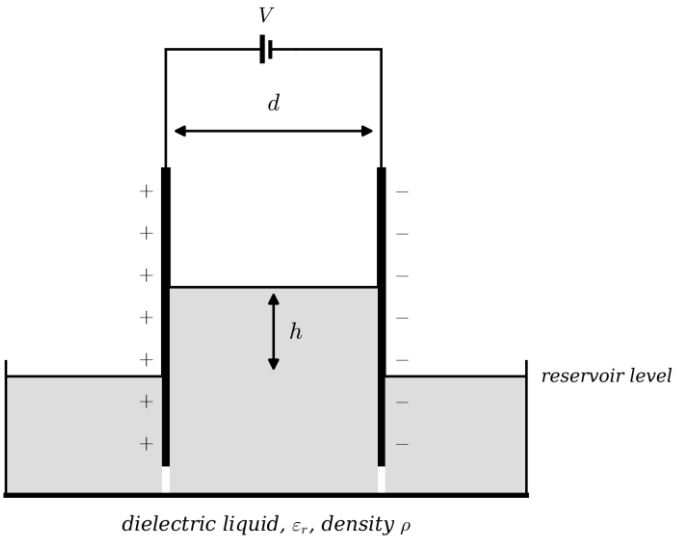


Figure: *A parallel-plate capacitor with vertical plates is dipped partway into a tank of dielectric liquid of dielectric constant ε_r and density ρ, with the plates connected to a battery of voltage V. The liquid rises inside the gap to a height h above the reservoir level outside the plates, and stays there in equilibrium.*

The height h can be found by an energy argument. Treat the gap as two capacitors in parallel, side by side: the dielectric-filled lower portion of height h and the air-filled upper portion of height $(\mathrm{H}-\mathrm{h})$, where H is the immersed plate height. With the plates connected to a battery of voltage V, the total capacitance is:

$$\mathrm{C(h)} = \frac{\varepsilon_0 \mathrm{L}}{\mathrm{d}}\big((\mathrm{H}-\mathrm{h}) + \varepsilon_\mathrm{r}\mathrm{h}\big) = \frac{\varepsilon_0 \mathrm{L}}{\mathrm{d}}(\mathrm{H} + (\varepsilon_\mathrm{r}-1)\mathrm{h})$$

Where L is the plate width and d is the gap. The voltage V is constant; the charge Q on the plates and the stored energy U change with h as the liquid rises.

By the energy method with a connected battery, the upward force on the dielectric liquid as it is drawn into the gap is:

$$\mathrm{F_{up}} = 1/2\,\mathrm{V}^2\frac{\mathrm{dC}}{\mathrm{dh}} = 1/2\,\mathrm{V}^2\frac{\varepsilon_0 \mathrm{L}(\varepsilon_\mathrm{r}-1)}{\mathrm{d}}$$

This force acts on the column of liquid above the reservoir level, lifting it. The lifted column has height h, width L, depth d (the gap), and density ρ, so its weight is:

$$\mathrm{W} = \rho \mathrm{Lhdg}$$

In equilibrium the electric lifting force equals the weight of the lifted column:

$$1/2\,\mathrm{V}^2\frac{\varepsilon_0 \mathrm{L}(\varepsilon_\mathrm{r}-1)}{\mathrm{d}} = \rho \mathrm{Lhdg}$$

Solving for h, with L cancelling from both sides:

$$\mathbf{h} = \frac{(\varepsilon_\mathrm{r}-1)\varepsilon_0 \mathrm{V}^2}{2\rho \mathrm{g}\mathrm{d}^2} = \frac{\boldsymbol{(\varepsilon_\mathrm{r}-1)\varepsilon_0 \mathrm{E}^2}}{\boldsymbol{2\rho g}},$$

where the right-hand form uses $\mathrm{E} = \mathrm{V}/\mathrm{d}$ for the electric field in the gap.

The rise is proportional to $(\varepsilon_\mathrm{r}-1)$, to the square of the field, and is inversely proportional to the density and to gravity. It does not depend on the plate width L, which is a good consistency check: a wider capacitor lifts a wider column with the same height. It also does not depend on the depth of immersion below the reservoir, provided the upper edge of the plates is above the rise level so the column has somewhere to go.

HOT Example 29

Two vertical parallel plates of width 5cm and separation 2mm are connected to a 5kV supply and partly immersed in a tank of transformer oil (dielectric constant 2.2, density $880\mathrm{kg/m^3}$). Find the height to which the oil rises inside the gap, above the reservoir level outside the plates. Take $\mathrm{g} = 9.81\mathrm{m/s^2}$.

Solution

The field between the plates is the voltage divided by the gap:

$$\mathrm{E} = \frac{\mathrm{V}}{\mathrm{d}} = \frac{5000\mathrm{V}}{2\times10^{-3}\mathrm{m}} = 2.5\times10^{6}\mathrm{V/m}$$

Apply the height formula derived above:

$$\mathrm{h} = \frac{(\varepsilon_\mathrm{r}-1)\varepsilon_0 \mathrm{E}^2}{2\rho \mathrm{g}} = \frac{(2.2-1)(8.854\times10^{-12}\mathrm{F/m})(2.5\times10^{6}\mathrm{V/m})^2}{2(880\mathrm{kg/m^3})(9.81\mathrm{m/s^2})} = 3.85\times10^{-3}\mathrm{m} = 3.85\mathrm{mm}$$

The transformer oil rises about 3.85mm inside the gap, above the level of the surrounding bath.

Making Sense of the Answer: *A rise of nearly four millimetres is easily visible to the eye and easily measurable with a ruler held against the outside of the capacitor. The plate width 5cm did not enter the answer, as the derivation showed; only the field, the dielectric constant, the density, and gravity.*

Think Like a Physicist: *The dimensions of the rise formula are worth checking. The numerator* $(\varepsilon_r - 1)\varepsilon_0 E^2$ *has units of energy per volume, since* $1/2\,\varepsilon_0 E^2$ *is the energy density of the field. The denominator* ρg *has units of weight per volume (kilogram per cubic metre times metre per second squared, which is newton per cubic metre, or equivalently joules per metre to the fourth). The ratio* J/m^3 *divided by* $N/m^3 = J/m^4$ *comes out cleanly in metres, as* h *must. Carrying units through every algebraic step, as we have done here, is the cheapest insurance against the kind of slip that drops a factor of* g *and gives an answer in inverse seconds.*

Conducting slab inserted in the gap

What happens if the slab inserted into the gap is a conductor rather than a dielectric? A perfect conductor cannot sustain an electric field inside itself; any field would drive its free electrons until they piled up on the surfaces and exactly cancelled the field in the bulk. So a conducting slab of thickness t in a parallel-plate gap has zero field inside it: all the voltage drop between the plates occurs in the air on either side of the slab.

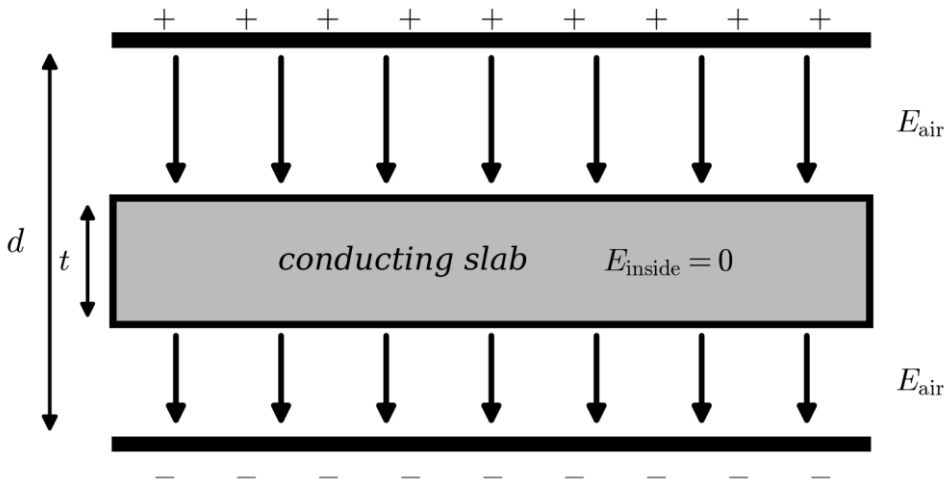


Figure: ***A conducting slab of thickness t inserted into a parallel-plate gap of width d.*** *The field inside the slab is zero. The field outside the slab, in the two air regions, is the full vacuum value* σ/ε_0*. The effective electrical gap is* $d - t$*, independent of where in the gap the slab sits.*

The effective electrical gap is therefore $\mathrm{d} - \mathrm{t}$, regardless of where the slab sits in the gap. The new capacitance is the vacuum capacitance with the reduced gap:

$$\mathrm{C} = \frac{\varepsilon_0 \mathrm{A}}{\mathrm{d} - \mathrm{t}}$$

A conducting slab can be thought of as the limit of a dielectric slab with $\varepsilon_\mathrm{r} \to \infty$. In the partial-fill formula, as ε_r becomes very large, the term $\mathrm{t}/\varepsilon_\mathrm{r}$ tends to zero, and the denominator reduces to $(\mathrm{d} - \mathrm{t})$, recovering the conducting-slab answer. The contrast between a dielectric and a conductor is exact in the limit: a dielectric reduces the field inside itself by a factor $1/\varepsilon_\mathrm{r}$; a conductor reduces it by a factor 0. Both raise the capacitance, but the conductor is the more efficient choice geometrically.

In practice a conducting slab is rarely used to fill part of the gap of a capacitor, because the slightest contact between the slab and either plate would short the capacitor. But the conducting-slab result is the cleanest baseline for partial-fill problems and the limiting case that calibrates the dielectric formulas.

HOT Example 30

An air-gap parallel-plate capacitor has plates separated by 5mm and capacitance 100pF. A conducting slab of thickness 3mm is inserted into the gap, without touching either plate. Find the new capacitance. Then find the thickness of a dielectric slab of $\varepsilon_\mathrm{r} = 4$ that would produce exactly the same new capacitance.

Solution

The original (air-only) capacitance, the gap, and so the product $\varepsilon_0 \mathrm{A}$, follow from the given:

$$\varepsilon_0 \mathrm{A} = \mathrm{C}_0 \mathrm{d} = (100 \times 10^{-12}\mathrm{F})(5 \times 10^{-3}\mathrm{m}) = 5 \times 10^{-13}\mathrm{Fm}$$

After the conducting slab is inserted, the effective gap is reduced to $\mathrm{d} - \mathrm{t}$:

$$d - t = 5\text{mm} - 3\text{mm} = 2\text{mm} = 2 \times 10^{-3}\text{m}$$

The new capacitance:

$$C = \frac{\varepsilon_0 A}{d - t} = \frac{5 \times 10^{-13}\text{Fm}}{2 \times 10^{-3}\text{m}} = 2.5 \times 10^{-10}\text{F} = 250\text{pF}$$

Or, using the ratio form, $C/C_0 = d/(d - t) = 5/2 = 2.5$, so $C = 2.5C_0 = 250\text{pF}$.

For the dielectric slab equivalent, set the partial-fill formula equal to 250pF and solve for t_d:

$$\frac{\varepsilon_0 A}{(d - t_d) + \frac{t_d}{\varepsilon_r}} = 250\text{pF}$$

The denominator must equal $\varepsilon_0 A/(250\text{pF}) = 2\text{mm}$:

$$(d - t_d) + \frac{t_d}{4} = 2\text{mm}$$

With $d = 5\text{mm}$:

$$5 - t_d + \frac{t_d}{4} = 2 \quad \Rightarrow \quad 5 - \frac{3t_d}{4} = 2 \quad \Rightarrow \quad t_d = 4\text{mm}$$

A dielectric of $\varepsilon_r = 4$ must be 4mm thick to produce the same 250pF as a 3mm conducting slab. The dielectric needs to be thicker to make up for not being a perfect conductor.

Making Sense of the Answer: *A 3mm conductor and a 4mm dielectric of $\varepsilon_r = 4$ are electrically equivalent for this geometry, both raising the capacitance from 100pF to 250pF. The conductor wins by a millimetre of thickness because it reduces the field inside itself by the factor $1/\infty$ rather than by the factor $1/4$. The factor-of-four dielectric makes up three quarters of the thickness it occupies, electrically speaking, but not all of it; the remaining quarter is "wasted" gap.*

Think Like a Physicist: *The conducting-slab result is the cleanest test of the partial-fill formula. If a student plugs $\varepsilon_r \to \infty$ into the partial-fill formula and does not recover the conducting-slab answer, the student has the formula wrong. Limiting cases are the cheapest, fastest, and most reliable formula checks available in physics. They cost almost no time and save a great deal of grief when an answer turns out to disagree with intuition.*

The dielectric is now thoroughly mapped: where it sits in the gap, how much of it sits there, in what configuration, with what battery condition, and at what temperature, all addressed and tabulated. But the capacitor and its insulators have, so far, all been in steady states. Nothing in this section has asked the question of how the charge gets onto the plates in the first place, or how long it takes. The next section watches the charge accumulate over time, through the resistance of the charging circuit, and exposes the second timescale that decides what a capacitor will actually do in a real device.

CHARGING AND DISCHARGING THROUGH A RESISTOR

The capacitor in the camera flash mentioned at the start of this chapter does its work in two distinct stages. Two small dry-cell batteries cannot deliver the 2000V that a flash bulb needs; what they can do is take their time, half a second or so, to push a small current through a high-resistance charging coil and into a $50\mu\text{F}$ capacitor, which sits and waits with its voltage rising slowly towards the supply value. When the shutter button is pressed, the capacitor is suddenly connected to the flash bulb through a near-zero resistance, and the stored energy comes out of the capacitor in under a millisecond, lighting the bulb so brightly that for one frame of motion the room is brighter than midday.

The capacitor here is acting as a bucket. The batteries fill the bucket slowly, through the narrow neck of the charging resistor. The flash bulb empties the bucket quickly, through the wide outlet of the low-resistance discharge path. The capacitor itself does not care which mode it is in; it simply holds whatever charge has been pushed onto it and lets whatever current the rest of the circuit allows flow into or out of it. What this section does is put numbers on the timing of both processes: *how long it takes to fill the bucket, and how long it takes to empty it.*

Charging through a resistor

The circuit for charging a capacitor through a resistor is shown below. A battery of EMF V drives current through a resistor R and into a capacitor C. Before the switch is closed, the capacitor is uncharged. The instant the switch is closed, current begins to flow, charge begins to build up on the plates, and the voltage across the capacitor begins to rise from zero.

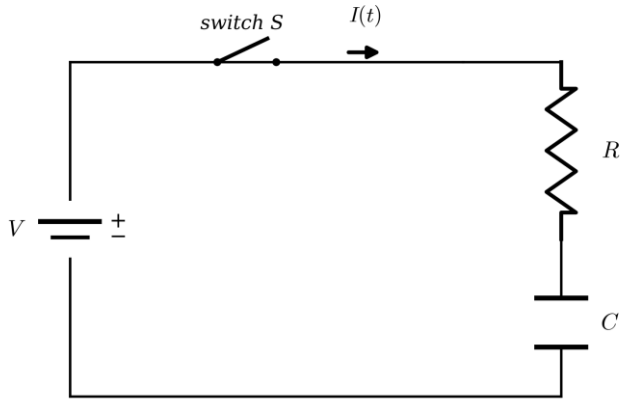


Figure: ***Charging circuit.*** *A battery of EMF V is connected in series with a switch S, a resistor R, and an initially uncharged capacitor C. When S is closed, current $I(t)$ flows and the capacitor charges towards a final voltage equal to the battery EMF.*

To find $\mathrm{Q(t)}$, the charge on the capacitor at time t after the switch is closed, apply **Kirchhoff's voltage law** around the loop. The sum of the potential drops around the closed loop must equal the EMF:

$$\mathrm{V = IR + \frac{Q}{C}}$$

The first term is the voltage drop across the resistor, the second is the voltage across the capacitor. The current I is the rate at which charge arrives on the capacitor plates: $\mathrm{I = \frac{dQ}{dt}}$. Substituting:

$$\mathrm{V = R\frac{dQ}{dt} + \frac{Q}{C}}$$

Rearrange to put Q and $\mathrm{\frac{dQ}{dt}}$ on one side:

$$\mathrm{R\frac{dQ}{dt} = V - \frac{Q}{C} = \frac{CV - Q}{C}}$$

Separate the variables, with Q and dQ on the left and t and dt on the right:

$$\mathrm{\frac{dQ}{CV - Q} = \frac{dt}{RC}}$$

Integrate both sides.

$$\mathrm{\int_0^Q \frac{dQ}{CV - Q} = \int_0^t \frac{dt}{RC}}$$

$$\mathrm{[-\ln(CV - Q)]_0^Q = \left[\frac{t}{RC}\right]_0^t}$$

$$\mathbf{\ln\frac{CV}{CV - Q} = \frac{t}{RC}}$$

Exponentiating both sides:

$$\mathrm{\frac{CV}{CV - Q} = e^{t/RC} \quad \Rightarrow \quad CV - Q = CVe^{-t/RC}}$$

Solving for $\mathrm{Q(t)}$:

$$\mathrm{Q = \mathbf{Q(t)} = \mathbf{CV\left(1 - e^{-t/RC}\right)}}$$

This is the charging law. The charge starts at zero when $\mathrm{t = 0}$, rises towards the asymptotic value:

$$\mathrm{Q_{max} = CV \text{ as } t \to \infty,}$$

and never actually reaches it; the approach is exponential. The current $\mathrm{I(t)}$ is the derivative of $\mathrm{Q(t)}$:

$$\mathbf{I(t) = \frac{dQ}{dt} = CV\left(\frac{1}{RC}e^{-t/RC}\right) = \frac{V}{R}e^{-t/RC}}$$

The current starts at the value $\mathrm{I_0 = \frac{V}{R}}$ at $\mathrm{t = 0}$, which is the value it would have if the capacitor were not there at all, and decays exponentially to zero as the capacitor fills and starts to oppose further charge flow.

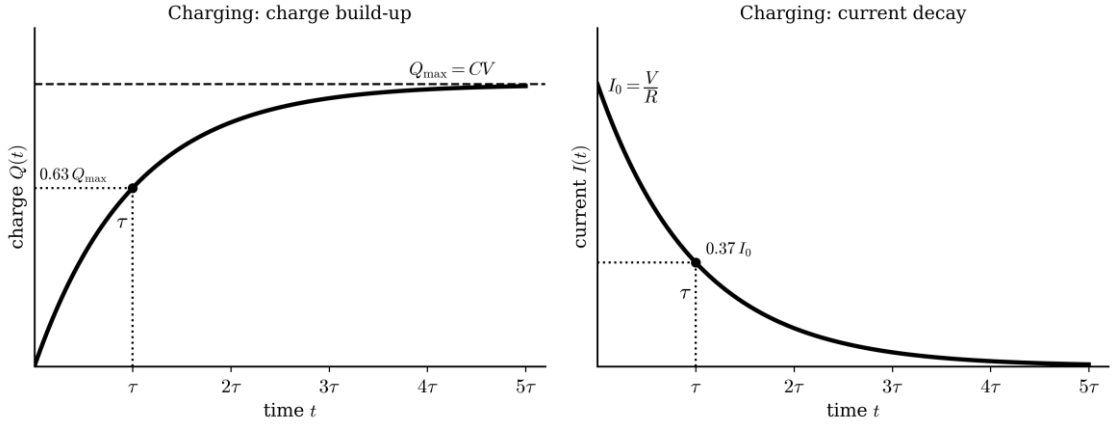


Figure: ***Charging curves.*** *Left: the charge $Q(t)$ rises from zero towards its asymptotic value $Q_{max} = CV$, reaching 63% of the asymptote at time $t = \tau$. Right: the current $I(t)$ starts at its maximum $I_0 = V/R$ and decays exponentially, falling to 37% of I_0 at $t = \tau$.*

The time constant

The combination RC that appears in both exponentials has units of time and is called the **time constant** of the circuit, written **τ**:

$$\boldsymbol{\tau = RC}$$

A quick dimensional check confirms that the product of resistance and capacitance is indeed a time:

$$[\mathrm{R}] \times [\mathrm{C}] = \frac{[\mathrm{V}]}{[\mathrm{i}]} \times \frac{[\mathrm{Q}]}{[\mathrm{V}]} = \frac{[\mathrm{Q}]}{[\mathrm{i}]} = \frac{[\mathrm{I}][\mathrm{T}]}{[\mathrm{I}]} = [\mathrm{T}]$$

The time constant is the natural timescale of the circuit. After one time constant, the charge has reached $1 - \mathrm{e}^{-1} = 0.632$ of its maximum value, often quoted as 63%. After two time constants, the charge has reached $1 - \mathrm{e}^{-2} = 0.865$ of its maximum. After five time constants, the charge has reached 0.993, within a per cent of its asymptotic value; for almost any practical purpose, the capacitor is fully charged after 5τ.

Concisely: **Time constant** of an electrical circuit *is the time taken for a changing quantity, such as charge, current, or voltage, to reach about 63% of its final value during growth, or to fall to about 37% of its initial value during decay.*

The resistance R does not change the final value $\mathbf{Q_{max} = CV}$ that the capacitor approaches; that is set by the battery EMF and the capacitance alone. What R changes is the speed at which $\mathrm{Q_{max}}$ is approached. *A large R gives a long time constant and a slow approach; a small R gives a short time constant and a fast approach.* This is the meaning of the bucket metaphor with which this section opened: *the bucket size is set by C and V, while the rate of filling is set by R.*

BINDER Example 31

A capacitor of capacitance $10\mu\mathrm{F}$ is to be charged from a 12V supply through a series resistor of $1.0\mathrm{M}\Omega$. Find the time constant of the circuit, and the time required for the capacitor to reach 99% of the supply voltage.

Solution

The time constant is the product of resistance and capacitance:

$$\tau = \mathrm{RC} = (1.0 \times 10^{6}\Omega)(10 \times 10^{-6}\mathrm{F}) = 10\mathrm{s}$$

The charge at time t obeys the charging law:

$$\frac{\mathrm{Q(t)}}{\mathrm{CV}} = \frac{\mathrm{Q(t)}}{\mathrm{Q_{max}}} = 1 - \mathrm{e}^{-\mathrm{t}/\tau}$$

Set this fraction equal to 0.99 and solve for t:

$$1 - \mathrm{e}^{-\mathrm{t}/\tau} = 0.99 \quad \Rightarrow \quad \mathrm{e}^{-\mathrm{t}/\tau} = 0.01$$

Take the natural logarithm of both sides:

$$-\frac{\mathrm{t}}{\tau} = \ln(0.01) = -\ln(100) = -4.605$$

So:

$$\mathrm{t} = 4.605\tau = 4.605 \times 10\mathrm{s} = 46.05\mathrm{s} \approx 46\mathrm{s}$$

The capacitor reaches 99% of the supply voltage in about 46s.

Making Sense of the Answer: *Forty-six seconds is just under five time constants, which is the rule-of-thumb point at which a capacitor is "fully" charged for engineering purposes. Each successive time constant gets the capacitor closer to Q_{max}, but by smaller and smaller increments: 63% after 1τ, 86% after 2τ, 95% after 3τ, 98% after 4τ, 99% after 4.6τ. The exponential approach means that "all the way" is never reached in finite time, but the time to reach any sensible fraction of the way is set by the time constant alone.*

Think Like a Physicist: *Whenever an exponential approach to a fixed value is asked about, the same algebra applies regardless of the physical setting. The same time-constant arithmetic will return in the next chapter when current grows through an inductor, and again in the chapter on radioactivity when a sample decays away. The trick "set the relevant ratio equal to the target fraction and take a logarithm" is one of the most reusable in physics. The harder skill is recognising which ratio belongs in the exponential: for a quantity that grows towards a limit, it is the deficit (the un-arrived-at remainder); for a quantity that decays towards zero, it is the quantity itself.*

HOT Example 32

A capacitor of capacitance $1.0\mu F$ is connected in series with a resistor of $10M\Omega$ and a $12.0V$ battery. Before the switch is closed, the capacitor is uncharged. At time $t = 46s$ after the switch is closed, find: (a) the fraction of the final charge that the capacitor has acquired; (b) the fraction of the initial current that is still flowing.

Solution

The time constant of the circuit is:

$$\tau = RC = (10 \times 10^6 \Omega)(1.0 \times 10^{-6} F) = 10s$$

At $t = 46s$, the dimensionless time variable is:

$$\frac{t}{\tau} = \frac{46s}{10s} = 4.6$$

The exponential factor at this time:

$$e^{-t/\tau} = e^{-4.6} = 0.01005$$

(a) Fraction of the final charge:

$$\frac{Q(t)}{Q_{max}} = 1 - e^{-t/\tau} = 1 - 0.01005 = 0.98995 \approx 0.99$$

The capacitor has acquired about 99% of its final charge.

(b) Fraction of the initial current still flowing:

$$\frac{I(t)}{I_0} = e^{-t/\tau} = 0.01005 \approx 0.01$$

The current has fallen to about 1% of its initial value.

Making Sense of the Answer: *The two results are mirror images. The charge has come 99% of the way; the current has fallen 99% of the way. Their sum is exactly one, every time, because $(1 - e^{-t/\tau}) + e^{-t/\tau} = 1$ is an identity. As the capacitor fills, the back-voltage it presents to the battery rises towards V and the net driving voltage falls towards zero, taking the current with it. The current and the deficit charge $(Q_{max} - Q)$ track each other exactly.*

Think Like a Physicist: *The complementary relation between $Q(t)/Q_{max}$ and $I(t)/I_0$ is more than an arithmetic curiosity. It says that the same exponential factor controls everything in the charging problem: the deficit charge, the instantaneous current, and through them the instantaneous power, the rate of energy storage, and the rate of resistor heating. Once $e^{-t/\tau}$ is computed for a given moment, every time-dependent quantity in the circuit can be read off without doing more exponentials. Spotting a single underlying scalar that controls many quantities is a habit worth developing; it cuts exam-paper arithmetic.*

Discharging through a resistor

Once a capacitor has been charged, it can be allowed to discharge through a resistor by disconnecting the battery and closing a switch that puts the capacitor directly across the resistor. The capacitor's own voltage now drives the current, which flows out of the positive plate, through the resistor, and back to the negative plate, gradually neutralising the charge.

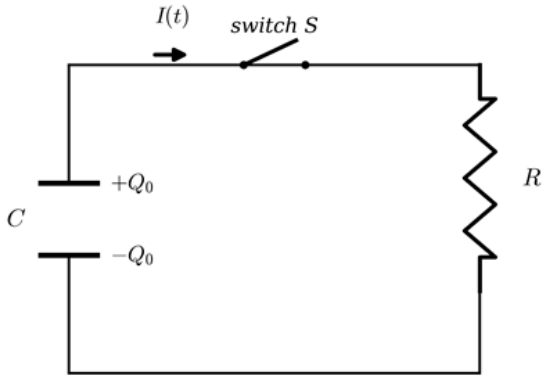


Figure: *Discharging circuit. A capacitor C, initially charged to Q_0, is connected through a switch S across a resistor R. When S is closed, the capacitor's stored charge drives a current $I(t)$ around the loop, decreasing Q to zero as $t \to \infty$.*

Kirchhoff's voltage law for the discharge loop says the voltage across the capacitor equals the voltage drop across the resistor:

$$\frac{Q}{C} = IR$$

The current flows out of the capacitor and so reduces Q, which means $I = -\frac{dQ}{dt}$ (the minus sign because Q is decreasing). Substituting:

$$\frac{Q}{C} = -R\frac{dQ}{dt}$$

Rearrange and separate variables:

$$\frac{dQ}{Q} = -\frac{dt}{RC}$$

Integrate both sides, from $Q = Q_0$ at $t = 0$ to Q at time t on the left, and from 0 to t on the right:

$$\int_{Q_0}^{Q}\frac{dQ}{Q} = \int_0^t -\frac{dt}{RC}$$

$$[\ln Q]_{Q_0}^{Q} = \left[-\frac{t}{RC}\right]_0^t$$

$$\boldsymbol{\ln\frac{Q}{Q_0} = -\frac{t}{RC}}$$

Exponentiating:

$$Q = \mathbf{Q(t) = Q_0 e^{-t/RC}}$$

The charge decays exponentially from Q_0 to zero, with the same time constant $\tau = RC$ as the charging case. The current is:

$$\mathbf{I(t) = -\frac{dQ}{dt} = \frac{Q_0}{RC}e^{-t/RC} = I_0 e^{-t/\tau}}$$

where $I_0 = \frac{Q_0}{RC} = \frac{V_0}{R}$ is the initial current, the same value one would get by treating the capacitor at $t = 0$ as a battery of voltage $V_0 = \frac{Q_0}{C}$ driving current through R. After one time constant, both the charge and the current have fallen to $1/e \approx 37\%$ of their initial values. After 5τ, both have fallen to less than 1% of their initial values; the capacitor is effectively discharged.

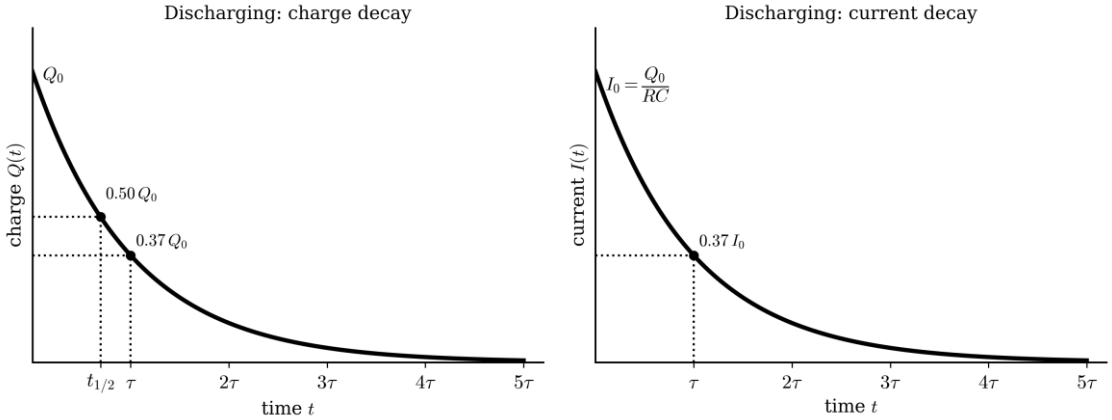


Figure: ***Discharging curves.*** *Left: the charge $Q(t)$ decays from Q_0 towards zero, falling to half its initial value at $t = t_{1/2} = \tau ln2$ and to 37% at $t = \tau$. Right: the current $I(t)$ decays from its initial value $I_0 = Q_0/(RC)$ along the same exponential.*

Half-discharge time

A useful single-number summary of how fast a capacitor discharges is the time required for the charge to fall to half its initial value. Set $Q(t) = \frac{Q_0}{2}$ in the discharge law:

$$\frac{Q_0}{2} = Q_0 e^{-t_{1/2}/\tau} \quad \Rightarrow \quad e^{-t_{1/2}/\tau} = \frac{1}{2}$$

Take the natural logarithm:

$$\frac{t_{1/2}}{\tau} = \ln 2 \approx 0.693$$

So:

$$\mathbf{t_{1/2} = \tau ln2 = 0.693\tau}$$

The half-discharge time is fixed at 0.693 time constants, independent of the initial charge. This is the same exponential half-life that appears for radioactive decay; the underlying mathematics of any exponential decay is the same regardless of whether the decaying quantity is the charge on a capacitor, the number of un-decayed nuclei in a radioactive sample, the concentration of a reactant in a first-order chemical reaction, the height of foam on a cup of cappuccino, or the height of froth on a glass of milk gradually collapsing with time.

The same half-time formula applies to charging, by symmetry: the deficit charge $(Q_{max} - Q)$ during charging halves every $\tau\ln 2$ seconds, because the deficit follows an exponential decay law:

$$(Q_{max} - Q) = Q_{max}e^{-t/\tau}$$

BINDER Example 33

A capacitor of $50\mu F$ is initially charged to $200V$ and then discharged through a $20k\Omega$ resistor. Find the time constant, the half-discharge time, and the time at which the voltage across the capacitor has fallen to $25V$.

Solution

The time constant:

$$\tau = RC = (20 \times 10^3\Omega)(50 \times 10^{-6}F) = 1.0s$$

The half-discharge time:

$$t_{1/2} = \tau\ln 2 = (1.0s)(0.693) = 0.693s$$

For the voltage to fall to $25V$ from $200V$, note that voltage and charge fall together (since $V = \frac{Q}{C}$ and C is fixed). The required fraction is $\frac{25V}{200V} = \frac{1}{8} = 0.125$.

$$\frac{V(t)}{V_0} = e^{-t/\tau} = 0.125$$

Take the natural logarithm:

$$-\frac{t}{\tau} = \ln(0.125) = -2.079$$

So:

$$t = 2.079\tau = (2.079)(1.0s) = 2.08s$$

Or by the half-time trick: $200V \to 100V \to 50V \to 25V$ is three halvings, which takes $3t_{1/2} = 3(0.693s) = 2.08s$. Both routes give the same answer.

Making Sense of the Answer: *The half-time route is the quicker check because the target voltage is exactly one-eighth of the initial value, which is* 2^{-3}. *Whenever the target is a power-of-two fraction of the initial value, the answer is an integer number of half-times. For target fractions that are not powers of two, the natural-logarithm route is the general method, but the half-time view still gives a quick estimate: at* $t_{1/2}$ *we are at* 50%, *at* $2t_{1/2}$ *at* 25%, *at* $3t_{1/2}$ *at* 12.5%, *and so on.*

Think Like a Physicist: *The exponential law gives the discharging capacitor a memoryless quality: from any moment onwards, the time required to fall to half of the present value is always* $t_{1/2} = \tau\ln 2$, *regardless of how much charge has already drained away. The first half-discharge takes* 0.693τ, *the second half-discharge from the halfway point also takes* 0.693τ, *and so on indefinitely. Memoryless decay is a hallmark of any first-order linear system, and recognising it lets the problem be solved by simple halving arguments instead of by the full exponential formula.*

REAL Example 34

A clinical defibrillator in a Dar es Salaam regional hospital uses a $100\mu F$ capacitor that is charged from a high-voltage supply and then discharged through the chest of a patient between two paddles. The chest tissue between the paddles has an electrical resistance of about $1.0k\Omega$. For one rescue, the capacitor is charged to $50V$ before the discharge. Find: (a) the initial discharge current; (b) the time constant of the discharge; (c) the half-discharge time; (d) the total energy delivered to the chest tissue; (e) the rise in temperature of a $1.0g$ packet of tissue between the paddles, if its specific heat capacity is the same as water, $4180J/(kgK)$.

Solution

(a) Initial discharge current:

At $t = 0$ the capacitor acts as a battery of voltage $V_0 = 50V$ across the resistance $R = 1.0k\Omega$:

$$I_0 = \frac{V_0}{R} = \frac{50V}{1.0 \times 10^3\Omega} = 0.050A = 50mA$$

(b) Time constant:

$$\tau = RC = (1.0 \times 10^3\Omega)(100 \times 10^{-6}F) = 0.10s = 100ms$$

(c) Half-discharge time:

$$t_{1/2} = \tau\ln 2 = (0.10s)(0.693) = 0.0693s \approx 69ms$$

(d) Total energy delivered:

All of the initial stored energy on the capacitor is dissipated in the resistor (there is nowhere else for it to go):

$$U_0 = 1/2\,CV_0^2 = 1/2\,(100 \times 10^{-6}F)(50V)^2 = 0.125J$$

(e) Temperature rise:

The heat absorbed by the $1.0g$ packet equals $mc\Delta T$, set equal to the energy delivered:

$$\Delta T = \frac{U_0}{mc} = \frac{0.125J}{(1.0 \times 10^{-3}kg)\big(4180J/(kgK)\big)} = 0.0299K \approx 30mK$$

The chest tissue between the paddles rises by about 30mK, or three hundredths of a kelvin, far too small to feel.

Making Sense of the Answer: *A temperature rise of three hundredths of a kelvin from a defibrillator-style discharge sounds suspiciously small, and it is. The reason is that the* $50V$ *used in this teaching example is far below the voltage of a real defibrillator. Clinical units charge the capacitor to between* $1500V$ *and* $3000V$*, where the stored energy is* 30^2 *to* 60^2 *times larger, so the energy delivered jumps from* $0.125J$ *to around* $100J$ *to* $400J$*, in line with the specifications of clinical equipment. The current also rises by the same factor of 30 to 60, from* $50mA$ *here to* 1.5 *to* 3 *amperes in the real machine, which is enough to deliver the shock that restarts a heart in fibrillation. The teaching numbers are scaled down purely for arithmetic simplicity; the physical formulas and time scales transfer directly to the real device.*

Think Like a Physicist: *A defibrillator is a deliberately fast discharge into a deliberately low resistance. The product* $RC = 100ms$ *here is short on a human timescale but exactly right for the heart: long enough to deliver substantial energy in one shock, short enough to leave the cardiac tissue unharmed between shocks. Designers of medical equipment do not pick R and C independently; they pick the product to match the biological time constant they need to interact with. Same trick is used in camera flashes (matched to the time the shutter is open), in radar pulse generators (matched to the round-trip travel time to the target), and in switching power supplies (matched to the switching frequency).* RC *is not an accident; it is a design knob.*

Energy bookkeeping during charging

A small but striking question hides inside the charging process. The battery does a known amount of work moving charge around the circuit. The capacitor stores a known amount of energy at the end. The resistor heats up by a known amount during the charging. The question is: *do these three numbers add up correctly, and what is the ratio between them?* The answer, the same in every charging problem regardless of the value of R, is surprising enough to be worth deriving carefully.

The total work done by the battery from the moment the switch is closed until the capacitor is fully charged is the time integral of the instantaneous power supplied by the battery, $\mathrm{P_{battery}} = \mathrm{VI(t)}$:

$$\mathrm{W_{battery}} = \int_0^\infty \mathrm{V\,I(t)dt} = \mathrm{V}\int_0^\infty \frac{\mathrm{V}}{\mathrm{R}}\mathrm{e^{-t/RC}dt}$$

The integral evaluates to RC, giving:

$$\mathrm{W_{battery}} = \frac{\mathrm{V^2}}{\mathrm{R}} \times \mathrm{RC} = \mathrm{CV^2}$$

The energy ultimately stored on the capacitor, from the standard formula:

$$\mathrm{U_{cap}} = 1/2\,\mathrm{CV^2}$$

The difference between battery work and stored energy must have been dissipated in the resistor as heat:

$$\mathrm{W_R} = \mathrm{W_{battery}} - \mathrm{U_{cap}} = \mathrm{CV^2} - 1/2\,\mathrm{CV^2} = 1/2\,\mathrm{CV^2}$$

So the battery does $\mathrm{CV^2}$ of work, exactly half of which ends up stored on the capacitor and exactly half of which is dissipated as heat in the resistor.

This 50-50 split is independent of the value of R. A small resistor charges the capacitor quickly with a large current; a large resistor charges slowly with a small current. Either way, exactly half of the battery's energy is dissipated as heat. This is the **half-energy theorem of capacitor charging:** *it is impossible to charge a capacitor from a battery without losing exactly half of the supplied energy as heat in the resistance of the circuit, no matter how small that resistance is.*

The result can be verified by direct integration of the resistor power. The instantaneous power dissipated in the resistor is $\mathrm{P_R(t)} = \mathrm{I^2(t)R}$, and the total heat is:

$$\mathrm{W_R} = \int_0^\infty \mathrm{I^2\,(t)Rdt} = \mathrm{R}\int_0^\infty \left(\frac{\mathrm{V}}{\mathrm{R}}\right)^2 \mathrm{e^{-2t/RC}dt} = \frac{\mathrm{V^2}}{\mathrm{R}}\int_0^\infty \mathrm{e^{-2t/RC}\,dt}$$

The integral evaluates to $\frac{\mathrm{RC}}{2}$:

$$\mathrm{W_R} = \frac{\mathrm{V^2}}{\mathrm{R}} \times \frac{\mathrm{RC}}{2} = 1/2\,\mathrm{CV^2}$$

The same $1/2\,\mathrm{CV^2}$, independent of R as advertised. Even in the limit $\mathrm{R} \to 0$, the current spikes to infinity for an infinitesimal time and the integral still evaluates to exactly $1/2\,\mathrm{CV^2}$. There is no resistance small enough to avoid the half-energy loss.

HOT Example 35

A $100\mu\text{F}$ capacitor is charged from a 50V battery through a 200Ω series resistor. Compute the total work done by the battery, the energy ultimately stored on the capacitor, and the heat dissipated in the resistor, and verify the half-energy theorem.

Solution

The total work done by the battery:

$$\mathrm{W_{battery}} = \mathrm{CV^2} = (100 \times 10^{-6}\mathrm{F})(50\mathrm{V})^2 = 0.25\mathrm{J}$$

The energy ultimately stored on the capacitor:

$$\mathrm{U_{cap}} = 1/2\,\mathrm{CV^2} = 1/2\,(0.25\mathrm{J}) = 0.125\mathrm{J}$$

The heat dissipated in the resistor, by the half-energy theorem:

$$W_R = 1/2\,CV^2 = 0.125J$$

Verification by direct integration of $P_R = I^2R$. The time constant is:

$$\tau = RC = (200\Omega)(100 \times 10^{-6}F) = 0.020s$$

The initial current:

$$I_0 = \frac{V}{R} = \frac{50V}{200\Omega} = 0.25A$$

The heat integral:

$$W_R = \int_0^\infty I_0^2\, e^{-2t/\tau} R dt = I_0^2 R \times \frac{\tau}{2} = (0.25A)^2(200\Omega)(0.010s) = 0.125J$$

The direct integration gives the same $0.125J$ that the half-energy theorem predicts. Energy bookkeeping closes: $W_{battery} = U_{cap} + W_R$, with $0.25J = 0.125J + 0.125J$, exactly.

Making Sense of the Answer: *The 200Ω resistor heated up to a noticeable degree during the $\approx 100ms$ of charging. The peak power dissipated in the resistor at $t = 0$ was $I_0^2R = (0.25)^2(200) = 12.5W$, a substantial wattage for a brief instant. By the time the capacitor was charged, the resistor had absorbed $0.125J$ of heat. For a typical wire-wound resistor of mass $1g$ and specific heat capacity similar to copper $(385J/(kg \cdot K))$, this is a temperature rise of $0.125/(10^{-3} \times 385) = 0.32K$, just detectable to the touch.*

Think Like a Physicist: *The R-independence of the half-energy theorem is the surprising part. Intuition might say: surely a smaller R wastes less energy, because I^2R would be smaller. The catch is that a smaller R also lets the current $I_0 = V/R$ be larger, by exactly the same factor. The two effects cancel out. Whenever a result turns out to be independent of a parameter that the algebra contained, there is usually a cancellation worth understanding, and the cancellation is often a hint of a deeper conservation principle. In this case, the principle is that work must be done to push charge against an increasing back-voltage, and that work is set by the geometry of the Q–V relationship alone.*

Heat dissipated during partial discharge

The last quantity to handle is the heat that comes out of a capacitor during a partial discharge: from time $t = 0$ to time $t = T$, how much energy has been dissipated in the resistor by the discharging capacitor? Once T is known the answer is a short calculation. The energy left on the capacitor at time T is:

$$U(T) = \frac{Q^2(T)}{2C} = \frac{Q_0^2}{2C} e^{-2T/\tau} = U_0 e^{-2T/\tau}$$

The energy dissipated is the difference between the initial and the remaining stored energies:

$$W_R(T) = U_0 - U(T) = U_0 - U_0 e^{-2T/\tau} = U_0\left(1 - e^{-2T/\tau}\right) = 1/2\,CV_0^2\left(1 - e^{-2T/\tau}\right)$$

Two limit checks confirm the formula. As $T \to \infty$, the exponential goes to zero and the heat reaches $1/2\,CV_0^2$, which is the entire initial stored energy as expected; the capacitor is then fully discharged. As $T \to 0$, the exponential goes to one and the heat is zero; nothing has discharged yet, so nothing has been dissipated. The formula passes both checks.

HOT Example 36

A capacitor of $200\mu F$ is charged to $100V$ and then discharged through a $1.0k\Omega$ resistor. Find the heat dissipated in the resistor during the first $100ms$ of the discharge.

Solution

The time constant:

$$\tau = RC = (1.0 \times 10^3\Omega)(200 \times 10^{-6}F) = 0.20s$$

The initial stored energy on the capacitor:

$$U_0 = 1/2\,CV_0^2 = 1/2\,(200 \times 10^{-6}F)(100V)^2 = 1.0J$$

At $T = 100ms = 0.10s$, the exponent in the partial-discharge formula is:

$$\frac{2T}{\tau} = \frac{2(0.10s)}{0.20s} = 1.0$$

The exponential factor:

$$e^{-2T/\tau} = e^{-1} = 0.3679$$

The heat dissipated from $t = 0$ to $t = T$:

$$W_R(T) = U_0\left(1 - e^{-2T/\tau}\right) = (1.0J)(1 - 0.3679) = 0.632J$$

About $0.63J$ of heat has been dissipated in the resistor during the first $100ms$ of the discharge, with the remaining $0.37J$ still stored on the capacitor.

Making Sense of the Answer: *In one time constant of charging or discharging, the capacitor moves from 100% to 37% of its initial state in voltage and charge, but it moves from 100% to $37^2\% = 14\%$ of its initial energy, because energy is*

quadratic in voltage. Equivalently, 63% of the charge has flowed but 86% of the energy has come out. The exponential decay is faster in energy than in charge by a factor of two in the exponent, because $U \propto Q^2$.

Think Like a Physicist: *The factor of two in the energy exponent traces back to the squaring in $U = Q^2/(2C)$. Any quantity that depends quadratically on an exponentially-decaying variable will decay exponentially with twice the rate of the original variable. The same idea explains why intensity in optics, power in alternating-current circuits, and kinetic energy in damped motion all decay twice as fast (in exponent) as the underlying field, current, or velocity. Squaring a decaying exponential is a habit-forming move in physics, and the doubling of the rate is the bookkeeping that follows.*

The charge dynamics of a capacitor are now complete. We have a formula for the charge as a function of time, for the current at any moment, for the time constant that sets the speed of approach, for the half-time at which the charge falls to half its value, and for the heat dissipated during any window of the charging or discharging process. The next section returns from the time-dependent picture to a steady-state geometric question: *what is the capacitance of an isolated conducting sphere, or of two concentric spheres, or of two coaxial cylinders?* These geometries close the chapter and prepare the ground for Chapter 13, where the capacitor reappears as a reactive element in alternating-current circuits.

CAPACITANCE OF CONDUCTORS

Every capacitor seen so far has had two clearly defined plates with an air or dielectric gap between them. But it is possible to assign a capacitance even to a single isolated conductor, with no second plate anywhere nearby. The trick is to recognise that the "second plate" is, in such cases, the surface of the universe itself, sitting at zero potential infinitely far away. With that convention, any charged conductor has a definite capacitance, set by its geometry and the medium around it. This section works out the capacitance of three important geometries (the isolated sphere, the concentric spherical capacitor, and the coaxial cable) and uses them to settle a few questions about how large or small a real capacitor might reasonably be.

1. Isolated conducting sphere

Consider a metal sphere of radius R held in vacuum, far from any other conductor, and carrying a total charge Q. From Chapter 10, the potential at the surface of such a sphere, measured against a reference at infinity, is:

$$V = \frac{Q}{4\pi\varepsilon_0 R}$$

The potential is the same everywhere on the surface, because the surface is an equipotential of a conductor in equilibrium. Inside the conductor, the potential is the same as at the surface, because there is no electric field inside a conductor in equilibrium; the potential is constant throughout the bulk and on the surface. Outside the conductor, the potential falls as $1/r$, exactly as if all the charge were concentrated at the centre.

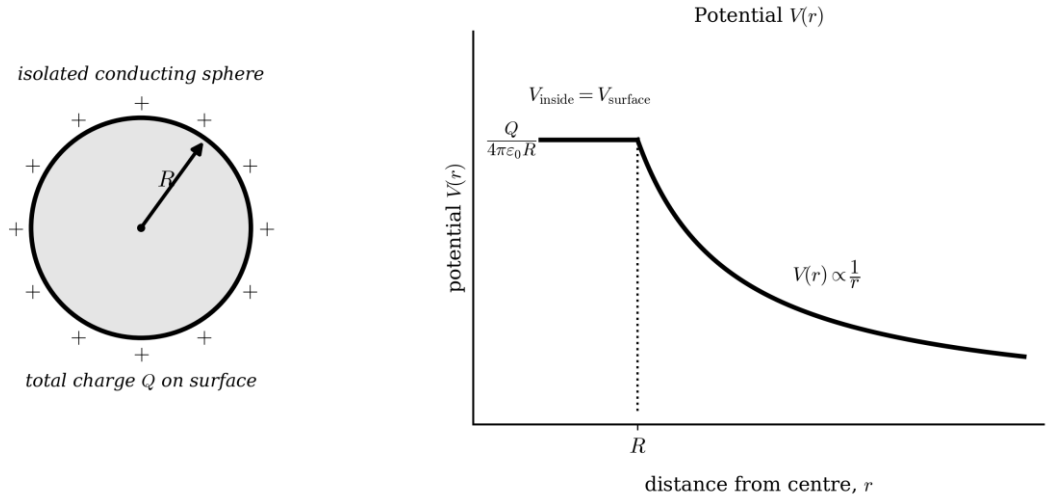


Figure: ***An isolated conducting sphere of radius R holds total charge Q on its surface.*** *The potential $V(r)$ is constant inside the conductor at the surface value $Q/(4\pi\varepsilon_0 R)$, and falls as $1/r$ outside.*

By the definition $C = \frac{Q}{V}$, the capacitance of the isolated sphere is:

$$\mathbf{C = \frac{Q}{V} = \frac{Q}{\frac{Q}{4\pi\varepsilon_0 R}} = 4\pi\varepsilon_0 R}$$

The capacitance depends only on the radius R and on ε_0. It is independent of the charge Q that the sphere happens to be carrying, and independent of any property of the conductor other than its size. A larger sphere can hold more charge at the same potential, and so has larger capacitance, by a factor exactly equal to the ratio of radii.

If the sphere sits in a dielectric medium of constant ε_r rather than in vacuum, the surface potential drops by the factor ε_r for the same charge, so the capacitance rises by the same factor:

$$\mathbf{C = 4\pi\varepsilon_r\varepsilon_0 R}$$

This is the same multiplicative rule that applied to the parallel-plate geometry: replace vacuum with a dielectric of constant ε_r and the capacitance is multiplied by ε_r regardless of geometry.

BINDER Example 37

A small isolated metal sphere of radius 12cm is held in vacuum, far from any other conductor. Find its capacitance.

Solution

Apply the isolated-sphere formula with $R = 0.12\text{m}$:

$$C = 4\pi\varepsilon_0 R = 4\pi(8.854 \times 10^{-12}\text{F/m})(0.12\text{m}) = 1.335 \times 10^{-11}\text{F} = 13.35\text{pF}$$

The capacitance of the isolated 12cm sphere is 13.35pF.

Making Sense of the Answer: *A sphere the size of a small football has a capacitance of only about 13pF. This is enormously less than a $1\mu F$ electrolytic capacitor the size of a fingernail, which is to say: making capacitance from an isolated charged conductor is wildly inefficient compared with making it from two close-spaced plates with a thin dielectric between them. The reason is geometric. The isolated-sphere "second plate" is at infinity, so the average distance between the plates is the radius itself, twelve centimetres in this case, and capacitance scales inversely with that gap. A parallel-plate capacitor with a $10\mu m$ dielectric gap shrinks the gap by a factor of ten thousand, gaining four orders of magnitude in capacitance per unit area at the cost of needing a second plate.*

Think Like a Physicist: *The remarkable fact about the isolated-sphere formula $C = 4\pi\varepsilon_0 R$ is that it has no V in it, no Q in it, no current in it; only the geometry and the constant of free space. Capacitance is a property of the conductor system, not of the state it happens to be in. The same lesson applies to all the geometries that follow: the formula always emerges from the geometry alone, with ε_0 and ε_r as the only material inputs. When a quantity depends only on geometry, it is robust against the kinds of accidents that can change the state of the system without changing what the system is.*

2. Concentric spherical capacitor

Now bring up a second conductor: a thin metallic shell of radius b, concentric with the first sphere, with the first sphere now of radius a, where $b > a$. The inner sphere carries charge $+Q$; the outer shell carries charge $-Q$, drawn from the earth to which it is connected. The region $a < r < b$ between the spheres holds an electric field, and integrating that field from a to b gives the voltage across the device.

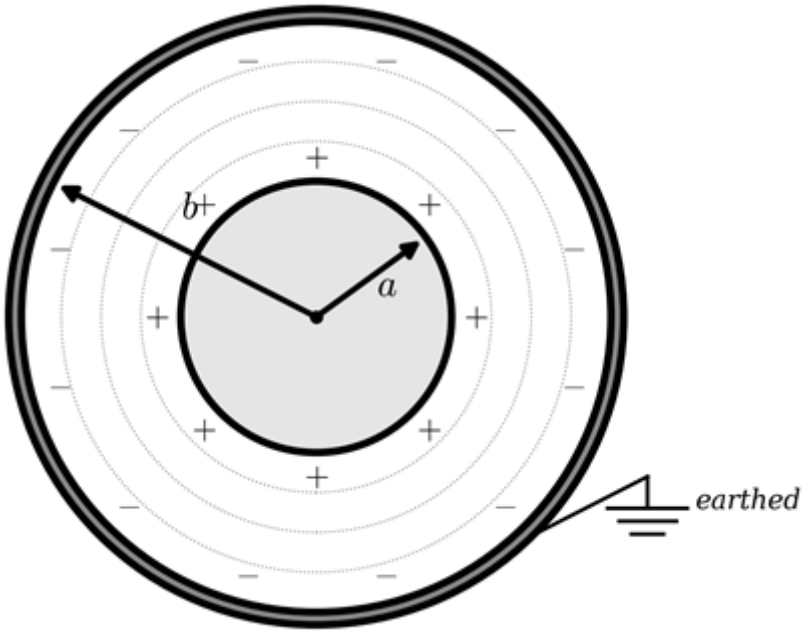


Figure: ***A concentric spherical capacitor.*** *The inner sphere has radius a and charge $+Q$; the outer shell has inner radius b and is earthed, drawing charge $-Q$ from the ground. The field in the annular gap $a < r < b$ falls as $1/r^2$. Outside the outer shell, the field is zero by Gauss's law.*

By Gauss's law applied to a spherical surface of radius r inside the gap, the field strength is:

$$E(r) = \frac{Q}{4\pi\varepsilon_0 r^2} \quad \text{for} \quad a < r < b$$

The potential difference between the inner and outer surfaces is the line integral of $E(r)$ from a to b:

$$V = \int_a^b E\,(r)dr = \frac{Q}{4\pi\varepsilon_0}\int_a^b \frac{dr}{r^2} = \frac{Q}{4\pi\varepsilon_0}\left[-\frac{1}{r}\right]_a^b = \frac{Q}{4\pi\varepsilon_0}\left(\frac{1}{a}-\frac{1}{b}\right)$$

Combine the two fractions:

$$V = \frac{Q}{4\pi\varepsilon_0}\left(\frac{b-a}{ab}\right)$$

The capacitance is $\frac{Q}{V}$:

$$\boldsymbol{C = \frac{Q}{V} = \frac{4\pi\varepsilon_0 ab}{b-a}}$$

Two limit checks confirm the formula sits between the two geometries we already understand.

Limit one: $b \to \infty$. As the outer shell is moved out to infinity, the formula reduces to $C \to 4\pi\varepsilon_0 a$, the isolated-sphere result for the inner sphere alone. The outer shell, far enough away, plays no role; the inner sphere "sees" only the universe at infinity.

Limit two: $b \to a$. As the outer shell is brought very close to the inner sphere, the gap $b-a$ becomes small while the product ab stays roughly a^2. The formula becomes $C \to 4\pi\varepsilon_0 a^2/(b-a)$. Identify $4\pi a^2$ as the surface area of the inner sphere, so this can be written $C \to \varepsilon_0 A/(b-a)$, exactly the parallel-plate formula with area $A = 4\pi a^2$ and gap $d = b-a$. The two-sphere geometry has collapsed to a parallel-plate geometry, as it should once the gap is much smaller than the radii.

Both limits work. The concentric spherical capacitor sits, formally, between the parallel-plate capacitor (gap much smaller than the curvature) and the isolated sphere (no outer plate at all), and it interpolates smoothly between them.

HOT Example 38

Treat the Earth as an isolated conducting sphere of radius $R_E = 6.4\times10^6\text{m}$ held in vacuum.

(a) Find the capacitance of the Earth, treated as an isolated sphere.

(b) Now treat the Earth together with the stratosphere (modelled as a thin conducting layer at altitude $h = 50\text{km}$ above the Earth's surface) as a concentric spherical capacitor. Find this capacitance.

(c) Compare the two values and explain which is larger and why.

Solution

(a) Apply the isolated-sphere formula with $R = 6.4\times10^6\text{m}$:

$$C_E = 4\pi\varepsilon_0 R_E = 4\pi(8.854\times10^{-12}\text{F/m})(6.4\times10^6\text{m}) = 712\times10^{-6}\text{F} = 712\mu\text{F}$$

The Earth, as an isolated sphere, has a capacitance of about $712\mu\text{F}$.

(b) Inner radius $a = R_E = 6.4\times10^6\text{m}$;

Outer radius $b = R_E + h = 6.4\times10^6\text{m} + 5.0\times10^4\text{m} = 6.45\times10^6\text{m}$. The gap is $b-a = 5.0\times10^4\text{m}$.

Apply the concentric-sphere formula:

$$C_{ES} = \frac{4\pi\varepsilon_0 ab}{b-a} = \frac{4\pi(8.854\times10^{-12}\text{F/m})(6.4\times10^6\text{m})(6.45\times10^6\text{m})}{5.0\times10^4\text{m}} = 0.0919\text{F} = 91.9\text{mF}$$

(c) The ratio of the two values:

$$\frac{C_{ES}}{C_E} = \frac{91.9\times10^{-3}\text{F}}{712\times10^{-6}\text{F}} = 129$$

The Earth-stratosphere capacitor is about 129 times larger than the isolated-Earth capacitance. The reason follows from the concentric-sphere formula: when the gap $b-a$ is small compared to ab, the capacitor is in the parallel-plate-like regime, and $C \approx \varepsilon_0 \cdot 4\pi a^2/(b-a)$. The stratosphere is only 50km above the surface, while the Earth's radius is 6400km, so the gap is just under 1% of the radius. In that thin-gap limit, the concentric formula gives a much larger capacitance than the isolated-sphere formula, by roughly the ratio of the radius to the gap, which is $6400\text{km}/50\text{km} = 128$, in excellent agreement with the numerical answer.

Making Sense of the Answer: *The Earth-and-stratosphere system has a capacitance of nearly* 0.1 *farads, an enormous value by any laboratory standard but unsurprising given the four-million-square-metre plate area. The atmospheric electrical circuit really does behave as a giant leaky capacitor: thunderstorms charge it from below, while a slow leakage current of about* $1{,}500A$ *worldwide closes the loop through the fair-weather atmosphere. The fact that this circuit can be sized using nothing more than the spherical-capacitor formula is a striking demonstration of how the simple geometries learned in this section apply to systems orders of magnitude larger than any device in the laboratory.*

Think Like a Physicist: *The crossover from the "isolated sphere" formula to the "parallel plate" formula happens whenever the outer shell is brought close enough to the inner sphere that the gap* $b-a$ *becomes much smaller than the radius* a*. This is exactly what happens to a beach ball when a metal sheet is brought up close to its surface; the ball, considered as an isolated conductor, has tiny capacitance, but bringing a near-by plate up to it raises the capacitance by the ratio of radius to gap. The lesson is that "isolated conductor" capacitance is a special case of two-conductor capacitance with the second conductor at infinity, and the second conductor matters whenever it is closer than the radius.*

REAL Example 39

Find the radius an isolated conducting sphere in vacuum would need to have a capacitance of exactly 1F. Compare with the radius of the Earth, and state whether such a capacitor would be a practical engineering object.

Solution

Rearrange the isolated-sphere formula to solve for the radius:

$$C = 4\pi\varepsilon_0 R \quad \Rightarrow \quad R = \frac{C}{4\pi\varepsilon_0}$$

Substitute $C = 1.0\text{F}$:

$$R = \frac{1.0\text{F}}{4\pi(8.854 \times 10^{-12}\text{F/m})} = 8.99 \times 10^9\text{m}$$

The radius needed is $8.99 \times 10^9\text{m}$, or roughly nine billion metres.

Compare with the Earth's radius:

$$\frac{R_{1F}}{R_E} = \frac{8.99 \times 10^9\text{m}}{6.4 \times 10^6\text{m}} \approx 1400$$

So a 1F isolated sphere would have to be about 1400 times the radius of the Earth. Such an object is not practical by any stretch of the imagination, and the one-farad rating, as a unit of capacitance, is fittingly enormous; capacitances of one farad and above exist in the laboratory only as electrolytic and supercapacitor constructions that achieve the rating through extreme plate area in a small volume, not through the use of any isolated sphere.

Making Sense of the Answer: *The one-farad isolated sphere would not even fit comfortably in the solar system. A sphere of radius $9 \times 10^9 m$ has a diameter of $1.8 \times 10^{10} m$, larger than the radius of Jupiter's orbit. The size of the farad as a unit of capacitance is therefore set by the inefficiency of the isolated-conductor geometry: a unit defined by $C = Q/V$ with Q in coulombs and V in volts has values, for any reasonable charged object, that are very small fractions of a farad. The microfarad, nanofarad, and picofarad are the units that actually appear on real components, and even the millifarad is on the large side for most circuits.*

Think Like a Physicist: *Whenever a unit comes out to "an impractically large value for any natural object", there is usually a story about which geometries can or cannot achieve the unit efficiently. A 1F capacitance is unreachable in a single isolated sphere of any natural size, but easily reached in an electrolytic capacitor the size of a thumb, because the parallel-plate geometry trades a large gap for a large area at small volume. The farad is large by the standards of isolated geometry but small by the standards of parallel-plate geometry, and the unit makes more sense when the latter is the reference.*

3. Coaxial (cylindrical) capacitor

The third geometry is the coaxial cable: two long concentric cylinders of radii a and b, with $b > a$ and length L, the inner cylinder carrying charge $+Q$ and the outer shell carrying $-Q$. This is the standard form of every transmission line: the inner conductor of a television antenna cable, the signal core of an oscilloscope lead, the centre conductor of any coaxial radio-frequency cable.

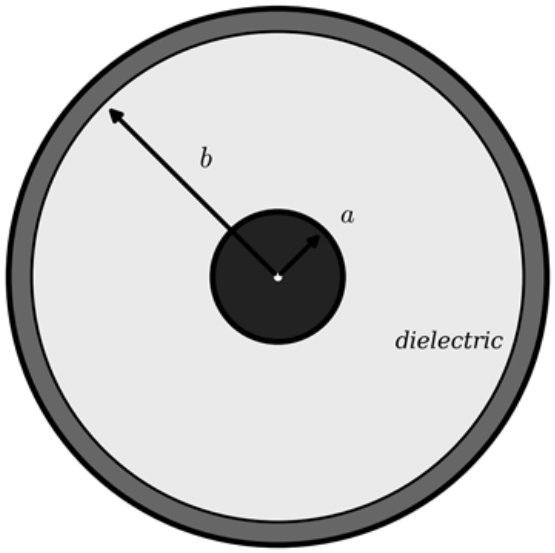


Figure: ***Cross-section of a coaxial cable.*** *The inner conductor of radius a carries charge $+Q$; the outer conductor of inner radius b carries the equal and opposite charge $-Q$. The dielectric between them, typically air, polyethylene, or foam, supports the field $E(r)$ that runs radially from inner to outer conductor.*

By Gauss's law applied to a cylindrical surface of radius r and length L, with the linear charge density $\lambda = \frac{Q}{L}$ on the inner conductor, the field in the gap is:

$$E(r) = \frac{\lambda}{2\pi\varepsilon_0 r} \quad \text{for} \quad a < r < b$$

The field falls as $1/r$, more slowly than the spherical $1/r^2$, because a cylinder spreads its flux over a circumference that grows linearly with r, while a sphere spreads its flux over a surface that grows quadratically.

Integrate the field from a to b to find the potential difference:

$$V = \int_a^b E\,(r)dr = \frac{\lambda}{2\pi\varepsilon_0}\int_a^b \frac{dr}{r} = \frac{\lambda}{2\pi\varepsilon_0}\ln\frac{b}{a}$$

Substitute $\lambda = \frac{Q}{L}$ to get V in terms of the total charge Q:

$$V = \frac{Q}{2\pi\varepsilon_0 L}\ln\frac{b}{a}$$

The capacitance follows from $C = Q/V$:

$$\boldsymbol{C = \frac{2\pi\varepsilon_0 L}{\ln(b/a)}}$$

This is the coaxial-cable capacitance formula. If the gap between the conductors is filled with a dielectric of constant ε_r, the formula is multiplied by ε_r in the usual way:

$$\boldsymbol{C = \frac{2\pi\varepsilon_r\varepsilon_0 L}{\ln(b/a)}}$$

Two features of the formula deserve comment. First, the capacitance depends on the logarithm of the radius ratio, not on the ratio itself; doubling the outer radius adds only $\ln 2 = 0.69$ to the denominator, modest compared with the strong dependence on length L. Second, the formula gives capacitance proportional to length, so the natural quoted figure for a cable is capacitance per unit length, C/L, in picofarads per metre. A typical RG-58 cable, with $\varepsilon_r \approx 2.3$, runs at about 100pF/m; a 30-metre run from the rooftop antenna to the receiver thus carries 3nF of stored capacitance that the signal has to drive through.

BINDER Example 40

A standard 50-ohm coaxial cable used in laboratory instrumentation has an inner conductor radius of 0.45mm and an outer conductor inner radius of 1.5mm. The gap between them is filled with polyethylene of dielectric constant $\varepsilon_r = 2.3$. Find the capacitance per metre of this cable.

Solution

Apply the coaxial formula divided by L:

$$\frac{C}{L} = \frac{2\pi\varepsilon_r\varepsilon_0}{\ln(b/a)}$$

$$\frac{C}{L} = \frac{2\pi(2.3)(8.854\times10^{-12}\text{F/m})}{\ln\left(\frac{1.5\text{mm}}{0.45\text{mm}}\right)} = 1.063\times10^{-10}\text{F/m} \approx 106\text{pF/m}$$

The cable has a capacitance per metre of about 106pF/m (in line with the manufacturer's specification of around 100pF/m for cables of this type).

Making Sense of the Answer: *A standard television cable from the roof to the set, running about* $20m$*, has a total cable capacitance of* $20\times106 = 2120pF \approx 2.1nF$*. This capacitance loads the antenna at high signal frequencies, where the cable's impedance drops to ohms and the antenna sees a heavy load. Cable specifications quote both the capacitance per metre and the characteristic impedance precisely because both numbers control the behaviour of the system at signal frequencies. The 50-ohm impedance and the* $100pF/m$ *capacitance are linked: changing one without changing the other would require changing the dielectric constant or the radius ratio, which would in turn change both.*

Think Like a Physicist: *The coaxial formula's* $\ln(b/a)$ *in the denominator reflects the radial geometry. In any flux problem where the field source is a line rather than a point, the field falls as* $1/r$ *rather than* $1/r^2$*, and the integral of* $1/r$ *over radius is a logarithm. The logarithm is a slowly varying function: doubling the radius ratio from 3 to 6 raises* $\ln(b/a)$ *from* 1.10 *to* 1.79*, a change of* 63%*, much less than the doubling of the geometry. Cable designers exploit this insensitivity to keep cable capacitance roughly constant even when the geometry has to flex for mechanical or manufacturing reasons.*

The three classical geometries (isolated sphere, concentric spheres, coaxial cylinders) now have their capacitances in closed form. They are the building blocks of every textbook capacitance calculation, and between them they cover the geometric range from "infinite gap" (isolated sphere) through "comparable gap and curvature" (concentric spheres) to "long cylinder with translation symmetry" (coaxial cable). The next section steps back from these analytical geometries to survey the practical applications of the capacitor: where it lives in everyday devices, what role it plays in each, and why the eight applications that close this chapter explain so much of the modern world.

APPLICATIONS OF CAPACITORS

Capacitors appear in almost every electronic device that exists, and in many that pass for non-electronic. The eight applications surveyed below cover the most important roles the capacitor plays in everyday equipment, from the camera flash that opened this chapter to the soil-moisture sensor on a maize plot. Each application is presented in two short paragraphs: what the capacitor is physically doing, and a few numbers or anchors that give the application its scale. None of the applications introduces new physics; everything here follows from the formulas already developed.

1. Flash photography

A camera flash unit charges a $50\mu F$ to $200\mu F$ capacitor from two AA batteries through a step-up oscillator to several hundred volts, taking about half a second. When the shutter is pressed, the capacitor is dumped through a xenon flash bulb of resistance a few ohms in under a millisecond, releasing all of its stored energy as light. The capacitor is the bucket from the opening of this chapter: filled slowly from a small source, emptied quickly into a low-resistance load.

A typical compact flash designed for $50J$ of light energy at a peak voltage of $300V$ requires a capacitor of $C = 2U/V^2 = 2(50)/(300)^2 F \approx 1.1mF$. The discharge time constant through a 10Ω flash bulb is $\tau = RC \approx 11ms$, with the light pulse lasting roughly $1ms$ to $5ms$ depending on the design. The high-pitched whine heard after pressing the flash button on an older camera is the oscillator stepping the battery voltage up to refill the bucket for the next shot.

2. Defibrillators

A clinical defibrillator stores around $200J$ of energy on a $100\mu F$ capacitor charged to about $2kV$, and discharges it through paddles pressed against a patient's chest in $5ms$ to $10ms$. The brief, intense current pulse depolarises the entire heart muscle simultaneously, allowing the sinoatrial node to restart a coordinated rhythm. The same capacitor formulas that apply to flash photography apply here, but at much higher voltage and energy.

The chest tissue between the paddles has a resistance of about $1k\Omega$, so the discharge time constant is $\tau = RC = 100ms$, intentionally short enough to deliver most of the energy in one decisive pulse and long enough not to be dangerously fast. This is the application worked through quantitatively in Example 34 of the previous section; the present paragraph simply explains why the numbers chosen there matter for the human heart.

3. DC blocking and AC coupling

A capacitor placed in series with a signal path blocks any DC component while passing AC, because a steady voltage produces no current through a capacitor in equilibrium but a varying voltage does. This is the standard way to feed an AC signal from one amplifier stage to the next without disturbing the DC bias of either; the capacitor passes the music or speech, the DC bias stays where the designer set it. The capacitor "couples" the AC signal across without "loading" the DC.

The capacitance is chosen so that the reactance $X_C = 1/(\omega C)$ is small compared with the load resistance at the lowest signal frequency of interest. For audio coupling from a $20Hz$ to $20kHz$ amplifier into a $10k\Omega$ stage, a $1\mu F$ capacitor gives $X_C = 1/\left(2\pi(20)(10^{-6})\right) \approx 8k\Omega$ at $20Hz$, comparable to the load and starting to roll off the bass. Chapter 13 takes up the reactance formula and what it does in alternating-current circuits in full.

4. Power-supply smoothing

A bridge rectifier converts $50Hz$ mains AC to a pulsating DC output at $100Hz$ (because both halves of the AC cycle now point the same way). A large capacitor placed in parallel with the load smooths these pulses by discharging slowly through the load between rectifier peaks, holding the load voltage almost constant. The capacitor is the reservoir; the rectifier refills it 100 times per second; the load drains it continuously.

The required capacitance grows with the load current: for a $1A$ load running off a $230V$ rectified supply with a ripple of 1% allowed, the capacitor must hold the voltage within $2.3V$ during the $10ms$ between peaks, giving $C = I\Delta t/\Delta V = (1)(0.010)/(2.3) \approx 4{,}400\mu F$. Large smoothing capacitors of $10{,}000\mu F$ at $400V$ working voltage are standard inside the cabinets of the Tanga substation and every other industrial DC supply in Tanzania.

5. Tuning circuits in AM radios

An inductor L placed in parallel with a variable capacitor C forms a resonant circuit that responds most strongly at the frequency $f_0 = 1/\left(2\pi\sqrt{LC}\right)$. A radio receiver uses this resonance to pick out one broadcast station from the dozens overlapping at its antenna: turning the tuning knob rotates one set of plates of the variable capacitor partly out of the gap between the fixed plates, changing the effective area A and so changing C, which shifts f_0 to the frequency of a different station.

A typical AM tuning capacitor varies from $20pF$ to $400pF$, paired with an inductor of around $250\mu H$, giving a tuning range from $f_0 \approx 510kHz$ to $f_0 \approx 2.3MHz$. The AM broadcast band that a Radio Free Africa receiver picks up runs from $540kHz$ to $1610kHz$, comfortably inside this range. The same trick is used in FM radios with smaller capacitors and inductors to reach the much higher FM frequencies.

6. Touch screens

Modern phone and tablet screens are tiled with a grid of tiny capacitors, each formed between two transparent conducting layers separated by a thin insulator. When a finger comes close to one of these capacitors, the finger acts as a dielectric of $\varepsilon_r \approx 50$ (because the human body is mostly salty water), raising the capacitance of just that one element by a measurable amount. The screen electronics scan the grid many times per second, detect which capacitor has changed, and compute the location of the touch from the position of the affected element.

Typical touch-screen elements have base capacitance around 5pF to 10pF and respond to a finger by changing by 0.5pF to 2pF, well above the noise floor of the detection electronics. A 5mm wide grid on a 10cm screen requires $20 \times 20 = 400$ capacitors and the same number of detection channels; modern phones have grids ten times finer than that. The capacitor formulas of this chapter, applied to elements only a few square millimetres in area, are what make capacitive touch interfaces commercially possible.

7. Electrostatic precipitators

A large industrial smoke duct can be cleaned by passing the flue gases between charged plates that act as the two plates of a capacitor on the kilowatt scale. The dust and ash particles in the gas pick up negative charge from a corona-discharge wire held at 30kV to 80kV, then drift across the duct under the field of the capacitor and stick to a positively charged collector plate. The collected dust is dislodged periodically by mechanical rapping and falls into hoppers below.

A typical precipitator at a cement works at Tanga Cement factory removes about 99.5% of the particulate matter from $100{,}000\,\text{m}^3/\text{h}$ of flue gas. The collector plates form a parallel-plate capacitor several metres on a side; the energy stored on the plates is modest, but the field strength of around 5kV/cm is enough to drive even the smallest charged dust grains across the gap in seconds. The same physics is used in coal-fired power stations and, on a much smaller scale, in domestic air purifiers.

8. Soil-moisture sensor for a Tanzanian maize plot

A pair of metal plates, each $5\text{cm} \times 5\text{cm}$ with a 2mm gap, is pushed vertically into the soil of a maize field outside Mbeya. The soil is the dielectric, and its dielectric constant rises sharply with water content: from about $\varepsilon_r = 4$ in dry soil to about $\varepsilon_r = 18$ in soil at field capacity. The capacitance of the sensor element therefore tracks soil moisture directly, and a small electronic module reads the capacitance and reports it as a percentage on a calibrated dial.

For the sensor described, the capacitance in dry soil is $C = \varepsilon_r \varepsilon_0 A/d = (4)(8.854 \times 10^{-12})(2.5 \times 10^{-3})/(2 \times 10^{-3}) \approx 44\text{pF}$, rising to about 200pF in fully saturated soil. The agricultural extension officer carries a hand-held unit and calls on each cooperative member once a week during the dry season to take a reading. Earlier in this chapter (Example 28) the temperature dependence of this same sensor was worked through in detail; the sensor is therefore the practical setting for both a capacitor formula and a dielectric-temperature correction.

The dead television in the prep room is now, retroactively, explained. The anode cap inside the picture tube was the storage plate of a high-voltage capacitor, several thousand volts above chassis ground; the glass envelope was the dielectric; and the chassis itself was the other plate. The slow leakage path that would eventually have drained the capacitor was through the (very small) conductivity of the glass and the surrounding air, with a time constant of weeks rather than seconds. This is why a technician shorts the anode cap to the chassis with a screwdriver before opening any television, even one unplugged for days. Disconnected, in this context, really is not the same as discharged.

Each of the eight applications above either fills its bucket slowly and empties it quickly (camera flash, defibrillator), or fills it fast and empties it slowly (smoothing capacitor), or holds a moderate charge that responds to small environmental changes (touch screen, soil-moisture sensor). Every one of these behaviours follows from a small set of formulas: the geometric capacitance $C = \varepsilon_0 A/d$, the energy storage $U = 1/2\,CV^2$, the time constant $\tau = RC$, and the dielectric multiplier $C = \varepsilon_r C_0$. The applications differ in geometry and in time scale but not in physics. Two chapters from now, the capacitor will reappear as a circuit element in alternating-current circuits, where the role of resistance is played by the reactance $X_C = 1/(\omega C)$; everything X_C depends on, the reader now knows how to compute.

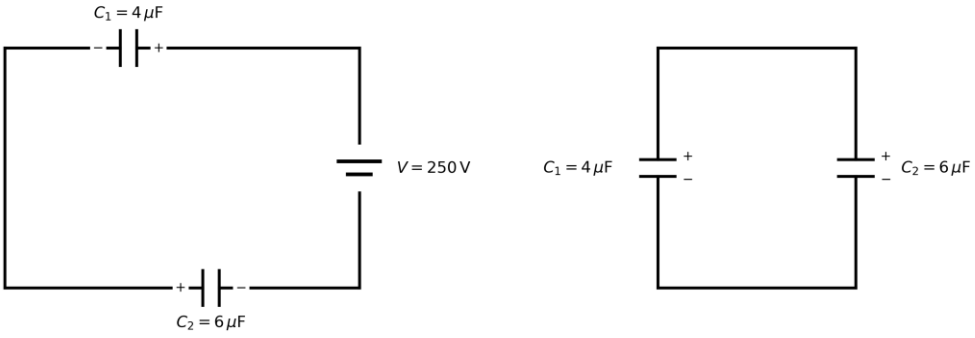
The chapter body is now complete. What follows is the Miscellaneous Worked Examples section, in which ten additional problems are worked through end to end.

MISCELLANEOUS WORKED EXAMPLES ON CAPACITANCE

Example 41

(a) State the conditions that must be met for a capacitor to be considered fully charged in a DC circuit, and explain briefly why these conditions are never strictly satisfied in finite time.

(b) Two capacitors connected as shown in the figure are first charged from the supply and are then disconnected and reconnected with like terminals joined. Calculate the new common potential difference and the new charge on each capacitor.



(a) Before disconnection (b) After reconnection

Solution

(a) A capacitor in a DC circuit is fully charged when:

(i) no current flows in the circuit, and

(ii) the potential difference across the capacitor equals the EMF of the source.

The charge during charging follows $Q(t) = CV\left(1 - e^{-t/RC}\right)$, which approaches CV only as $t \to \infty$. For any finite time t, the factor $e^{-t/RC}$ is strictly positive, so $Q(t)$ is strictly less than CV and a small current is still flowing.

(b) **Series** combination across the $250V$ supply. The equivalent capacitance is:

$$C_s = \frac{C_1 C_2}{C_1 + C_2} = \frac{(4\mu F)(6\mu F)}{4\mu F + 6\mu F} = 2.4\mu F$$

In a series chain each capacitor carries the same charge, equal to the charge supplied by the battery:

$$Q = C_s V = (2.4\mu F)(250V) = 600\mu C$$

Reconnection with **like** terminals joined. After disconnection each capacitor still carries $600\mu C$, and when the positive plates are joined to one common rail and the negative plates to another, the total charge on the joined positive rail is:

$$Q_{total} = Q_1 + Q_2 = 600\mu C + 600\mu C = 1200\mu C$$

The capacitors now sit in **parallel,** so the combined capacitance of the parallel pair is:

$$C_p = C_1 + C_2 = 4\mu F + 6\mu F = 10\mu F$$

The common potential difference across the parallel combination is the total charge divided by the combined capacitance:

$$V' = \frac{Q_{total}}{C_p} = \frac{1200\mu C}{10\mu F} = 120V$$

The new charges on each capacitor are therefore:

$$Q_1' = C_1 V' = (4\mu F)(120V) = 480\mu C$$
$$Q_2' = C_2 V' = (6\mu F)(120V) = 720\mu C$$

The new common potential difference is $120V$. The new charges are $480\mu C$ on the $4\mu F$ capacitor and $720\mu C$ on the $6\mu F$ capacitor.

Example 42

(a) Commercial paper capacitors are often manufactured by rolling two long strips of metal foil with the paper dielectric folded as a zig-zag between them, rather than laying the foil and the paper out flat. State two reasons why the zig-zag construction is preferred over flat-plate construction.

(b) Two capacitors of $2.0\mu F$ and $4.0\mu F$ are connected to form a potential divider, with the $4.0\mu F$ capacitor observed to have $5.0V$ across it. The input supply voltage is $15V$. Find the two possible values of energy stored in the $2.0\mu F$ capacitor, and explain in physical terms why two values are possible.

Solution

(a) (i) A long zig-zag of metal foil packs a much greater effective plate area into a small physical volume than a flat-plate stack of the same outer dimensions, so the capacitance per unit volume is much higher; this is the dominant reason.

(ii) The zig-zag pattern breaks the single long dielectric interface of a flat-plate stack into many shorter segments, so a single pinhole or microscopic defect in the paper is less likely to span the gap between the foils and trigger premature breakdown.

(b) The two possibilities arise because the polarity of the $5.0V$ across the $4.0\mu F$ capacitor is not specified relative to the $15V$ supply. The two possible values of the potential difference across the $2.0\mu F$ capacitor are:

$$V_{2,(a)} = 15V - 5V = 10V \quad \text{or} \quad V_{2,(b)} = 15V + 5V = 20V$$

The corresponding energies stored in the $2.0\mu F$ capacitor are:

$$U_{(a)} = \frac{1}{2}CV^2 = \frac{1}{2}(2.0\times10^{-6}\text{F})(10\text{V})^2 = 1.0\times10^{-4}\text{J} = 100\mu\text{J}$$
$$U_{(b)} = \frac{1}{2}CV^2 = \frac{1}{2}(2.0\times10^{-6}\text{F})(20\text{V})^2 = 4.0\times10^{-4}\text{J} = 400\mu\text{J}$$

The two possible energy values are $100\mu\text{J}$ and $400\mu\text{J}$.

Physically, if the 5V drop across the $4.0\mu\text{F}$ capacitor opposes the supply, the $2.0\mu\text{F}$ must supply the remaining 10V to make up the loop; if the same 5V aids the supply, the $2.0\mu\text{F}$ must drop the full 20V to satisfy Kirchhoff's voltage rule.

Example 43

(a) Explain briefly why the half-life of a discharging capacitor depends only on the product RC, and not on the initial charge stored or the voltage that originally charged the capacitor.

(b) A capacitor of capacitance $8.0\mu\text{F}$ is charged through negligible resistance to a potential difference of 18V, and is then discharged through a resistor of resistance $2.5\text{M}\Omega$. Find (i) the time constant of the discharge circuit; (ii) the charge remaining on the capacitor 8.0s after the discharge begins; (iii) the potential difference across the capacitor at the same instant; (iv) the current in the resistor at the same instant.

Solution

(a) The discharge of a capacitor through a resistor follows $Q(t) = Q_0\,e^{-t/RC}$, where Q_0 is the initial charge. The half-life is the time at which $Q(t)/Q_0 = 1/2$, which gives $e^{-t_{1/2}/RC} = 1/2$, and hence $t_{1/2} = RC\ln 2$. The initial charge Q_0 appears identically on both sides of the ratio and cancels, and the original charging voltage enters Q_0 only through $Q_0 = CV_0$ and similarly cancels. The half-life therefore depends only on the product RC.

(b) (i) The time constant is the product of resistance and capacitance:

$$\tau = R\,C = (2.5\times10^{6}\Omega)(8.0\times10^{-6}\text{F}) = 20\text{s}$$

(ii) The initial charge is:

$$Q_0 = C\,V_0 = (8.0\mu\text{F})(18\text{V}) = 144\mu\text{C}$$

The charge remaining at $t = 8.0\text{s}$ is:

$$Q(8) = Q_0\,e^{-t/\tau} = (144\mu\text{C})\,e^{-8.0/20} = (144\mu\text{C})(0.6703) = 96.5\mu\text{C}$$

(iii) The potential difference across the capacitor at the same instant follows from $V = Q/C$:

$$V(8) = \frac{Q(8)}{C} = \frac{96.5\mu\text{C}}{8.0\mu\text{F}} = 12.1\text{V}$$

(iv) The current in the resistor at the same instant follows from Ohm's law applied to the resistor:

$$I(8) = \frac{V(8)}{R} = \frac{12.1\text{V}}{2.5\times10^{6}\Omega} = 4.83\mu\text{A}$$

Example 44

(a) Show by direct argument that the force per unit area between the plates of a parallel-plate capacitor is numerically equal to the energy density of the field between the plates.

(b) A parallel-plate capacitor has plate area 0.05m^2 and plate separation 1.0mm, and is charged to a potential difference of 200V. (i) Find the force of attraction between the plates while the capacitor remains connected to the battery. (ii) The plates are then drawn apart slowly, with the battery still connected, until the separation is 2.0mm. Find the new force. (iii) Find the work done by the external agent in increasing the separation.

Solution

(a) The total electric field between the plates of a parallel-plate capacitor has magnitude $E = \sigma/\varepsilon_0$, where $\sigma = Q/A$ is the surface charge density on either plate. This total field is the sum of the field $\sigma/(2\varepsilon_0)$ contributed by each plate separately, and a plate cannot exert a force on itself, so only the field due to the **other** plate acts on the charge σA of the first. The force per unit area on one plate is therefore the surface charge density times the field due to the other plate alone:

$$\frac{F}{A} = \sigma\times\frac{\sigma}{2\varepsilon_0} = \frac{\sigma^2}{2\varepsilon_0}$$

Substituting $\sigma = \varepsilon_0 E$ recasts this in terms of the electric field:

$$\frac{F}{A} = \frac{(\varepsilon_0 E)^2}{2\varepsilon_0} = \frac{1}{2}\varepsilon_0 E^2$$

The right-hand side is exactly the energy density $u = \frac{1}{2}\varepsilon_0 E^2$ of the field between the plates.

Hence $F/A = u$: force per unit area equals energy per unit volume.

(b) Plate area $A = 0.05\text{m}^2$, initial separation $d_0 = 1.0\text{mm} = 1.0\times10^{-3}\text{m}$,
final separation $d_f = 2.0\text{mm} = 2.0\times10^{-3}\text{m}$, supply voltage $V = 200\text{V}$.

(i) The field while the capacitor sits at the initial separation is:

$$E_0 = \frac{V}{d_0} = \frac{200V}{1.0 \times 10^{-3}m} = 2.0 \times 10^5 V/m$$

The force of attraction follows from $F = \frac{1}{2}\varepsilon_0 E^2 A$:

$$F_0 = \frac{1}{2}(8.854 \times 10^{-12} F/m)(2.0 \times 10^5 V/m)^2 (0.05 m^2) = 8.85 \times 10^{-3} N$$

(ii) With the battery still connected, the voltage is held at $200V$. The new field at the doubled separation is:

$$E_f = \frac{V}{d_f} = \frac{200V}{2.0 \times 10^{-3}m} = 1.0 \times 10^5 V/m$$

The new force, with E halved, is therefore:

$$F_f = \frac{1}{2}(8.854 \times 10^{-12} F/m)(1.0 \times 10^5 V/m)^2 (0.05 m^2) = 2.21 \times 10^{-3} N$$

(iii) Because the battery holds V constant, the attractive force depends on the separation as $F(d) = \frac{\varepsilon_0 A V^2}{2 d^2}$, and the external agent must do positive work against this force in pulling the plates apart. The work is:

$$W_{ext} = \int_{d_0}^{d_f} F(d)\, dd = \int_{d_0}^{d_f} \frac{\varepsilon_0 A V^2}{2d^2} dd = \frac{\varepsilon_0 A V^2}{2} \int_{d_0}^{d_f} \frac{dd}{d^2} = \frac{\varepsilon_0 A V^2}{2} \left(\frac{1}{d_0} - \frac{1}{d_f}\right)$$

Substituting the numerical values:

$$W_{ext} = \frac{(8.854 \times 10^{-12})(0.05)(200)^2}{2} \left(\frac{1}{1.0 \times 10^{-3}} - \frac{1}{2.0 \times 10^{-3}}\right) = 4.43 \times 10^{-6} J$$

Example 45

(a) For a capacitor that must operate at high voltage without dielectric breakdown, explain why the dielectric strength of the material matters more than its dielectric constant. Use this to explain why mica capacitors are preferred over electrolytic capacitors for radio-frequency tuning circuits.

(b) A parallel-plate capacitor of capacitance $5.0nF$ is to be designed to operate at a working voltage of $600V$. The available dielectric is mica with relative permittivity $\varepsilon_r = 6$ and dielectric strength $E_b = 100MV/m$. A safety factor of 2 on the field is required. Find (i) the minimum plate separation that prevents breakdown at the working voltage; (ii) the required plate area; (iii) state one practical reason why the design might still fail in service even with the safety factor applied.

Solution

(a) The voltage at which a capacitor breaks down is $V_{max} = E_b d$, where E_b is the dielectric strength of the filling material and d is the plate separation. A higher relative permittivity ε_r increases the capacitance and hence the charge stored at a given voltage, but it does not change V_{max}. For a high-voltage capacitor the limiting design factor is therefore the dielectric strength, not the dielectric constant. Mica has a high dielectric strength of about $100MV/m$ and very low dielectric loss at radio frequencies, so it can hold off the peak voltages of an RF resonant circuit without breakdown and without dissipating the signal as heat. An electrolytic capacitor has a low dielectric strength (set by the thin oxide layer) and a high dielectric loss at radio frequencies; the same RF voltage would either break down its oxide or convert most of the signal energy into heat. Mica is therefore preferred.

(b) (i) With a safety factor of 2 applied to the field, the design field at the working voltage is:

$$E_{design} = \frac{E_b}{2} = 50MV/m = 5.0 \times 10^7 V/m$$

The minimum plate separation is:

$$d_{min} = \frac{V_{working}}{E_{design}} = \frac{600V}{5.0 \times 10^7 V/m} = 1.2 \times 10^{-5} m = 12\mu m$$

(ii) The required plate area follows from $C = \varepsilon_r \varepsilon_0 A/d$:

$$A = \frac{C d_{min}}{\varepsilon_r \varepsilon_0} = \frac{(5.0 \times 10^{-9} F)(1.2 \times 10^{-5} m)}{(6)(8.854 \times 10^{-12} F/m)} = 1.13 \times 10^{-3} m^2$$

(iii) The data-sheet figure of $100MV/m$ is the dielectric strength of clean, dry, defect-free mica. In service the mica sheet may carry microscopic voids, edge moisture absorption, or surface contamination, any of which can reduce the local breakdown field well below the data-sheet value and trigger premature failure of the device. A safety factor on the average field does not protect against a local field concentration at a defect that the bulk calculation does not see.

Example 46

(a) Earthing one plate of a charged parallel-plate capacitor does not change the capacitor's capacitance. Explain why.

(b) A parallel-plate capacitor of plate area $250cm^2$ with plates $2.0mm$ apart is fully charged from a $150V$ battery. With the battery still connected, the plates are then drawn apart slowly to a new separation of $4.0mm$. (i) State, with brief justification, how each of the field E, the charge Q on each plate, and the stored energy U changes during the process. (ii)

Compute E, Q, and U before and after the separation, and account in one sentence for where the lost capacitor energy went.

Solution

(a) Capacitance depends only on the geometry of the two-conductor system (the plate area A and the gap d) and on the dielectric between the plates. Earthing one plate fixes its potential at zero relative to earth, but it does not change A, d, or the dielectric, so the ratio $C = Q/V$ between the plates is unchanged.

(b) (i) With the battery connected the potential difference V is held fixed at 150V. The field $E = V/d$ halves as d doubles. The capacitance $C = \varepsilon_0 A/d$ halves as d doubles. The charge $Q = CV$ halves because C halves and V is fixed. The stored energy $U = 1/2\, CV^2$ halves for the same reason.

(ii) $A = 250\text{cm}^2 = 0.025\text{m}^2$, $d_0 = 2.0\text{mm} = 2.0 \times 10^{-3}\text{m}$, $d_f = 4.0\text{mm} = 4.0 \times 10^{-3}\text{m}$, $V = 150\text{V}$.

Initial configuration:

$$E_0 = \frac{V}{d_0} = \frac{150\text{V}}{2.0 \times 10^{-3}\text{m}} = 7.5 \times 10^4\text{V/m}$$

$$C_0 = \frac{\varepsilon_0 A}{d_0} = \frac{(8.854 \times 10^{-12}\text{F/m})(0.025\text{m}^2)}{2.0 \times 10^{-3}\text{m}} = 1.107 \times 10^{-10}\text{F} = 110.7\text{pF}$$

$$Q_0 = C_0 V = (110.7\text{pF})(150\text{V}) = 16.6\text{nC}$$

$$U_0 = \frac{1}{2} C_0 V^2 = \frac{1}{2}(110.7 \times 10^{-12}\text{F})(150\text{V})^2 = 1.25\mu\text{J}$$

Final configuration, with $d_f = 2d_0$:

$$E_f = \frac{V}{d_f} = 3.75 \times 10^4\text{V/m}, \qquad C_f = \frac{C_0}{2} = 55.4\text{pF}$$

$$Q_f = C_f V = 8.30\text{nC}, \qquad U_f = \frac{1}{2} C_f V^2 = 0.623\mu\text{J}$$

The capacitor returned exactly half of its initial stored energy ($0.623\mu\text{J}$) to the battery as charge flowed off the plates, and the external agent did an additional $0.623\mu\text{J}$ of mechanical work in pulling the plates apart against the residual attraction; both contributions ended up in the battery, which therefore gained a net $1.25\mu\text{J}$ in total.

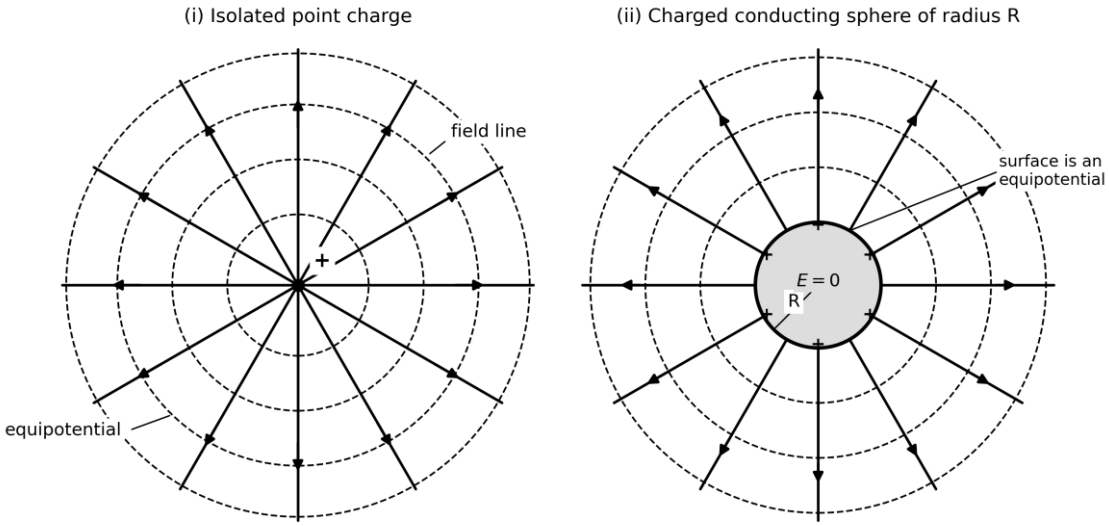
Example 47

(a) Sketch the electric field lines and the equipotential surfaces for (i) an isolated positive point charge in vacuum, and (ii) a charged conducting sphere of radius R, both inside and outside the sphere. State in each case why the equipotential surfaces are everywhere perpendicular to the field lines.

(b) A parallel-plate capacitor has plate area 4m^2 and the gap between the plates is completely filled by three layers of dielectric material stacked perpendicular to the plates. The three layers have relative permittivities 3, 6, and 9, and thicknesses 1.0mm, 3.0mm, and 0.6mm respectively. Calculate the capacitance of the capacitor.

Solution

(a) The sketches are shown in the figure below.



The equipotential surfaces are everywhere perpendicular to the field lines because V is constant on an equipotential surface; any small displacement $d\vec{s}$ along the surface gives $dV = -\vec{E} \cdot d\vec{s} = 0$, so $\vec{E}$ has no component along the surface and must lie

wholly perpendicular to it. For the point charge, the equipotentials are concentric spheres centred on the charge, and the field lines are radial straight lines piercing them at right angles. For the conducting sphere, the entire interior is at one potential and the field inside is zero; on the outside, the field and equipotentials look identical to those of a point charge of the same total charge placed at the centre, with the sphere's own surface as the innermost equipotential.

(b) The three dielectric layers, all with the same plate area A and stacked perpendicular to the plates, act as three capacitors in series. The combined capacitance follows from the parallel-plate formula applied to a stack:

$$C = \frac{\varepsilon_0 A}{\frac{d_1}{\varepsilon_{r1}} + \frac{d_2}{\varepsilon_{r2}} + \frac{d_3}{\varepsilon_{r3}}}$$

$$C = \frac{(8.854 \times 10^{-12}\text{F/m})(4\text{m}^2)}{\left(\frac{1.0}{3} + \frac{3.0}{6} + \frac{0.6}{9}\right) \times 10^{-3}\text{m}} = 3.94 \times 10^{-8}\text{F} = 39.4\text{nF}$$

The capacitance of the three-layer arrangement is 39.4nF.

Example 48

(a) Liquid dielectrics whose molecules are polar, such as water and propylene carbonate, have relative permittivities that decrease as the temperature is raised. Explain the molecular reason for this behaviour, and contrast it with the behaviour of non-polar dielectrics such as paraffin oil.

(b) A capacitor of unknown capacitance C is fully charged from a 200V battery. The battery is then disconnected and the capacitor is discharged through a small coil of resistance wire embedded in a thermally insulated block of specific heat capacity $250\text{J kg}^{-1}\text{K}^{-1}$ and mass 0.10kg. The temperature of the block rises by 0.40K as a result of the discharge. Calculate C.

Solution

(a) A polar liquid owes its high relative permittivity to the alignment of its permanent molecular dipoles with the applied field; the larger the aligned fraction of dipoles, the larger the induced opposing field inside the medium, and the larger the relative permittivity. Thermal motion opposes this alignment: the higher the temperature, the more vigorous the random reorientation of molecules and the smaller the fraction of dipoles that remain aligned at any instant, so ε_r falls. A non-polar liquid like paraffin oil has no permanent molecular dipoles; its (much smaller) ε_r comes from induced dipoles whose magnitudes depend only on the local field strength and not on the temperature, so ε_r is essentially independent of T.

(b) The energy initially stored in the capacitor is delivered entirely to the resistance-wire coil during the discharge, and the coil delivers it as heat to the thermally insulated block. The thermal energy gained by the block is therefore equal to the initial stored energy of the capacitor:

$$\frac{1}{2} C V^2 = m c \Delta T$$

Substituting the given quantities:

$$\frac{1}{2} C (200\text{V})^2 = (0.10\text{kg})(250\text{J kg}^{-1}\text{K}^{-1})(0.40\text{K})$$

$$C = \frac{2 \times 10\text{J}}{40000\text{V}^2} = 5.0 \times 10^{-4}\text{F} = 500\mu\text{F}$$

The capacitance is $500\mu\text{F}$.

Example 49

(a) An air capacitor of negligible self-capacitance is connected across an electrostatic voltmeter and the system is charged to 80V. A glass slab of dielectric constant 5 is then inserted to fill the gap of the capacitor completely while the voltmeter remains connected. The voltmeter reading is observed to fall to 24V. Explain why the voltmeter reading does not fall to the value $(80/5)\text{V} = 16\text{V}$ expected from the dielectric constant alone, and identify the source of the discrepancy.

(b) Two capacitors of capacitance C_1 and $C_2 = 12\mu\text{F}$ are connected in series and charged through negligible resistance from a battery. The battery is then disconnected and the series combination is discharged through a resistor of resistance $R = 20\text{M}\Omega$. The measured time constant of the discharge is $\tau = 60\text{s}$. Find C_1.

Solution

(a) The electrostatic voltmeter has its own (small but non-zero) capacitance C_V, which sits in parallel with the air capacitor C_0, and the initial charge on the system is $Q = (C_0 + C_V)(80\text{V})$, not $C_0(80\text{V})$. When the glass slab is inserted, the air capacitor becomes $5C_0$, but the voltmeter is not affected, so the new parallel combination has capacitance $5C_0 + C_V$. Charge is conserved during insertion, so the new reading is $V' = (C_0 + C_V)(80\text{V})/(5C_0 + C_V)$, which exceeds $80\text{V}/5 = 16\text{V}$ because the voltmeter contributes a capacitance that the dielectric does not enhance. The discrepancy comes from the parasitic capacitance of the voltmeter itself.

(b) The two capacitors in series behave electrically as a single capacitor of capacitance:

$$C_s = \frac{C_1 C_2}{C_1 + C_2}$$

The time constant for discharging through R is $\tau = R C_s$, so:

$$C_s = \frac{\tau}{R} = \frac{60s}{20 \times 10^6 \Omega} = 3.0 \times 10^{-6} F = 3.0 \mu F$$

Substituting back into the series formula and solving for C_1:

$$\frac{1}{C_1} = \frac{1}{C_s} - \frac{1}{C_2} = \frac{1}{3.0\mu F} - \frac{1}{12\mu F} = \frac{4-1}{12\mu F} = \frac{1}{4.0\mu F}$$

$$C_1 = 4.0\mu F$$

The unknown capacitance is $4.0\mu F$.

Example 50

(a) Explain how a projected-capacitance touch screen on a mobile phone detects the location of a finger.

(b) Two identical small conducting spheres, each of mass 5.0g and carrying a charge of $60\mu C$, are suspended from a common point by light insulating threads of equal length 30cm. In air, the threads hang at 36° to each other. The arrangement is now fully immersed in a transparent insulating liquid of density $800 kg\, m^{-3}$, and the angle between the threads is observed to remain unchanged. The density of the sphere material is $2400 kg\, m^{-3}$. Determine the relative permittivity of the liquid.

Solution

(a) A projected-capacitance touch screen is built as a fine grid of tiny capacitor cells, each formed by two transparent electrodes separated by a thin dielectric. When a finger approaches a cell, the finger sits in the cell's stray field and acts locally as a high-permittivity dielectric (the human body is mostly water, ε_r of order 50), replacing part of the air (ε_r close to 1) above the cell. That cell's capacitance therefore rises by tens of per cent. A controller chip scans every cell hundreds of times a second, identifies the cell whose capacitance has changed, and reports the grid coordinates of that cell as the touch location.

(b) Consider the following figure:

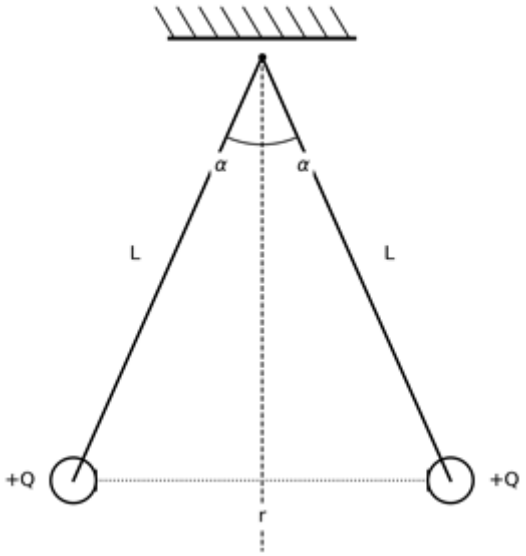


Figure: ***Geometry of the suspension in air***. *Each thread of length L makes the half-angle α with the vertical, the spheres are separated by r, and each sphere carries charge $+Q$.*

Set the half-angle to α, so that $2\alpha = 36°$ and $\alpha = 18°$. The geometry is the same in air and in the liquid; only the forces acting on each sphere change.

In air, three forces act on each sphere: the tension T along the thread, the weight mg downward, and the horizontal Coulomb repulsion F_C from the other sphere. Resolving the equilibrium horizontally and vertically:

$$T\sin\alpha = F_C, \qquad T\cos\alpha = mg \quad \Rightarrow \quad F_C = mg\tan\alpha$$

The Coulomb force in air is given by:

$$F_C = \frac{kQ^2}{r^2} \qquad (\text{with} \qquad r = 2L\sin\alpha)$$

In the liquid the angle is unchanged, so the separation r between the spheres is unchanged. But two things have changed:

- The Coulomb force is reduced by the relative permittivity of the liquid (from Chapter 10, the force between two charges in a dielectric medium is:

$$F = \frac{kQ_1Q_2}{\varepsilon_r r^2} = \frac{kQ^2}{\varepsilon_r r^2} = \frac{F_C}{\varepsilon_r}$$

- The effective weight of each sphere is reduced by the buoyant upthrust of the liquid. The buoyant force on each sphere is $V_{sphere}\,\rho_L\,g$, where $V_{sphere} = m/\rho_s$. The apparent weight in the liquid is therefore:

$$m'g = mg - \frac{m}{\rho_s}\,\rho_L\,g = mg\left(1 - \frac{\rho_L}{\rho_s}\right)$$

Equilibrium in the liquid gives:

$$\frac{F_C}{\varepsilon_r} = mg\left(1 - \frac{\rho_L}{\rho_s}\right)\tan\alpha$$

Dividing this by the equilibrium equation in air, $F_C = mg\tan\alpha$, makes the unknowns m, Q, L, and α all cancel cleanly, leaving the densities alone:

$$\frac{1}{\varepsilon_r} = 1 - \frac{\rho_L}{\rho_s} \quad \Rightarrow \quad \varepsilon_r = \frac{\rho_s}{\rho_s - \rho_L}$$

Substituting the given densities:

$$\varepsilon_r = \frac{2400\text{kg m}^{-3}}{2400\text{kg m}^{-3} - 800\text{kg m}^{-3}} = \frac{2400}{1600} = 1.50$$

The relative permittivity of the liquid is 1.50.

Notice that the answer is independent of the charge, the mass, the thread length, and the angle of suspension: only the density ratio of the sphere material to the liquid enters.

Ten miscellaneous worked examples now lie behind us. The Digging Deeper Exercise now follows.

DIGGING DEEPER EXERCISE 11

EXERCISE 11A: BINDER QUESTIONS

Question 1

A capacitor stores charge, yet the net charge on a fully charged capacitor is zero. Reconcile these two statements.

Question 2

Identify four factors that determine the capacitance of a parallel-plate capacitor, and indicate in each case whether the capacitance is directly or inversely proportional to that factor.

Question 3

Two capacitors of capacitances C_1 and C_2 are connected first in series and then in parallel across the same battery. In which configuration is the total charge delivered by the battery greater? Justify your answer.

Question 4

On the same pair of axes, sketch the variation of charge Q with time t for (i) a capacitor charging through a resistor from a constant-voltage source, and (ii) the same capacitor discharging through the same resistor. Mark the time constant τ on both curves.

Question 5

Give two reasons why mica is preferred over paper as the dielectric for a high-frequency tuning capacitor.

Question 6

Explain how an electrostatic precipitator cleans smoke from a chimney.

Question 7

Distinguish between the dielectric constant of a material and its dielectric strength. Explain why water is rarely used as a capacitor dielectric in spite of its high dielectric constant.

Question 8

Show by dimensional analysis that the product RC has the units of time.

Question 9

When a capacitor is being charged through a resistor by a constant-voltage source, the current in the resistor decreases steadily with time. Explain why the current does not stop suddenly at the instant the capacitor reaches its maximum charge.

Question 10

In an RC charging circuit, the series resistor R affects the time taken to reach the steady-state charge on the capacitor, but not the value of that steady-state charge. Identify the quantity that does determine the steady-state charge, and explain physically why R does not.

Question 11

Sketch the electric field lines and the equipotential surfaces between the two oppositely charged plates of a parallel-plate capacitor, including the regions near the edges of the plates.

Question 12

When the two plates of a parallel-plate capacitor have unequal areas, do they carry charges of equal magnitude after the capacitor has been connected to a battery? Justify your answer by appealing to conservation of charge.

EXERCISE 11B: REAL QUESTIONS

Question 13

The owner of an old film camera complains that the flash takes longer to recharge between shots as the batteries weaken. Explain why this is so.

Question 14

A simple washing-machine timer uses an RC circuit to switch off the motor after a fixed interval. If the time constant required is 10s and the resistor available is $100\text{k}\Omega$, estimate the required capacitance.

Question 15

Turning the tuning knob of an old AM radio rotates the moving plates of a variable capacitor partly out of the fixed plates. Explain how this changes the station received.

Question 16

A piezoelectric lighter produces a spark when the trigger is pressed. The spark lasts only a fraction of a second even if the trigger is held down. By treating the piezoelectric element as a small capacitor, explain why the spark is not sustained.

Question 17

Electrolytic capacitors are marked with a polarity and the manufacturer warns the user not to connect them with the polarity reversed. Explain in physical terms what happens if the polarity is reversed.

Question 18

A small torch designed for use in emergency power cuts contains a capacitor in parallel with the bulb, which continues to glow faintly for about 30s after the batteries die. Given that the bulb has resistance 100Ω and operates at 3V, estimate the capacitance of the torch's capacitor.

Question 19

A parallel-plate soil-moisture sensor placed in a maize plot near Mbeya reads 4nF in dry conditions and 15nF a day after heavy rain. Explain why the reading changes by such a large factor.

Question 20

When the switch of a torch with a capacitor across the bulb is opened, the bulb continues to glow for nearly a second before going dark. Explain this observation.

Question 21

Lightning between a cloud and the ground is often described as the discharge of a giant capacitor. Identify the two plates of this capacitor and the dielectric between them, and explain why the breakdown occurs only when the electric field exceeds approximately $3 \times 10^{6}\text{V/m}$ for dry air.

EXERCISE 11C: HOT QUESTIONS

Question 22

Three capacitors of capacitance $2\mu\text{F}$, $4\mu\text{F}$, and $6\mu\text{F}$ are connected first in series and then in parallel across the same 240V DC supply. In each case find the equivalent capacitance, the total charge delivered by the supply, and the total energy stored.

Question 23

Three $4\mu\text{F}$ capacitors are arranged in a network with two of them in parallel and the combination in series with the third. The network is connected across a 120V DC supply. Find (i) the equivalent capacitance; (ii) the charge on each capacitor; (iii) the potential difference across each capacitor; (iv) the total energy stored.

Question 24

A $2\mu\text{F}$ capacitor is connected to a 20V supply through a $3\text{M}\Omega$ series resistor and an open switch. Find the time taken after the switch is closed for the energy stored in the capacitor to reach 50 per cent of its maximum final value.

Question 25

A $50\mu\text{F}$ capacitor in series with a $4\text{M}\Omega$ resistor discharges through the resistor. Find the half-life of the discharge, and the number of half-lives needed for the charge to fall below 1 per cent of its initial value.

Question 26

Two square plates of side 10cm separated by 1.0mm carry charges of magnitude 8nC each, with opposite signs. Find the force of attraction between the plates and compare it with the weight of 1g of water.

Question 27

A $2\mu\text{F}$ capacitor carries a charge of 10mC and is discharged through a $5\text{M}\Omega$ resistor. Find (i) the charge remaining on the capacitor after 2s; (ii) the heat dissipated in the resistor during those 2s.

Question 28

A parallel-plate capacitor of plate area A and gap d has the gap filled by two dielectric slabs of relative permittivities $\text{K}_1 = 2$ and $\text{K}_2 = 4$, each of thickness $\text{d}/2$, stacked perpendicular to the plates. Show that the equivalent capacitance is:

$$\text{C}_{\text{eq}} = \frac{2\varepsilon_0 \text{A}}{\text{d}}\left(\frac{\text{K}_1\text{K}_2}{\text{K}_1 + \text{K}_2}\right)$$

Question 29

A capacitor of capacitance $20\mu\text{F}$ charged to 500V is short-circuited through a coil of resistance 2Ω. Find (i) the peak current immediately after the short-circuit; (ii) the time constant of the discharge; (iii) the total energy dissipated as heat in the coil.

Question 30

Two concentric thin metallic spherical shells have radii 5cm and 10cm. The inner shell carries a charge $+Q$ and the outer shell is earthed. Show that the system behaves as a spherical capacitor, find its capacitance, and compare it with the capacitance of an isolated sphere of radius 5cm.

Question 31

A coaxial cable consists of an inner cylindrical conductor of radius 1mm and an outer cylindrical conductor of radius 5mm, separated by an insulating dielectric of relative permittivity 2. The cable is 100m long. Find its capacitance.

Question 32

A parallel-plate capacitor with air between its plates has a capacitance of 8pF. Find the new capacitance if the distance between the plates is reduced by half and the space between them is filled with a substance of dielectric constant 6.

Question 33

A Van de Graaff generator has its shell electrode held at a potential of $2.5 \times 10^6\text{V}$. The dielectric strength of the gas surrounding the electrode is $5 \times 10^7\text{V/m}$. Find the minimum radius of the spherical shell that will prevent breakdown.

Question 34

A $100\mu\text{F}$ capacitor is charged by a 12V battery and then disconnected. It is then discharged through a thin copper wire of mass 2g and specific heat capacity $380\text{J kg}^{-1}\text{ K}^{-1}$. Assuming all the stored energy goes into heating the wire, find the rise in temperature of the wire.

Question 35

A simple parallel-plate soil-moisture sensor for a Tanzanian maize plot has plates of side 5cm and gap 2mm. Calculate (i) the capacitance in air; (ii) the capacitance in dry soil of relative permittivity 4; (iii) the capacitance in saturated soil of relative permittivity 18; (iv) suggest a simple circuit arrangement that converts the capacitance reading into a meter deflection a farmer can read in the field.

Question 36

In an RC discharging circuit, the time taken for the charge on the capacitor to fall to half its initial value is found to be 0.50s. Determine (i) the time constant of the circuit; (ii) the time taken for the charge on the capacitor to fall from 8.0mC to 1.0mC.

Question 37

A parallel-plate capacitor with plates of area 250cm^2 and separation 6.0mm is charged to a potential difference of 240V from a battery, after which the battery is disconnected. A slab of relative permittivity 5 and thickness 4.0mm is then inserted into the gap, leaving the remaining 2.0mm as air. Find (i) the new charge on each plate; (ii) the new potential difference between the plates; (iii) the new energy stored.

Question 38

A capacitor of capacitance $3.0\mu\text{F}$ is charged from one source to a potential difference of 80V, and a separate capacitor of capacitance $5.0\mu\text{F}$ is charged from another source to a potential difference of 200V. Both sources are then disconnected, and the capacitors are wired together with their like-polarity plates joined. Find (i) the energy stored in each capacitor before the connection; (ii) the total energy stored after the connection; (iii) the energy lost in making the connection, and account in one sentence for the loss.

Chapter 12

CURRENT ELECTRICITY

INTRODUCTION

Late one cloudy afternoon Mr. Akilikubwa reached for the switch on the laboratory wall. The fluorescent tubes overhead came on instantly, with the small click and pink flicker that Miono students had come to know. Kipanga, sitting at the front bench, was paying attention this time.

Mr. Akilikubwa: *Kipanga, you closed the switch and the room lit up at once. How fast did the electricity travel?*

Kipanga (confident)**:** *At the speed of light, sir. Three hundred thousand kilometres per second.*

Mr. Akilikubwa: *The signal did, yes. But the electrons themselves crawl. In the copper wire above your head, the electrons are drifting toward the bulbs at about half a millimetre every second. Slower than an ant carrying a leaf.*

Kipanga (suddenly less confident)**:** *But the bulb lit up immediately, sir.*

Kipute (working it out before Mr. Akilikubwa can)**:** *Of course it did. Every electron in the wire was already there before the switch closed. When the switch closed, every electron started pushing the next one along at the same instant. The bulb lit up because the wire was already full of electrons, not because any one electron had reached the bulb.*

Mr. Akilikubwa: *Kipute has it. The wire is like a long queue of people passing buckets. No person walks the length of the queue, but a bucket arrives at the end almost the moment the first bucket is handed in. The signal travels at the speed of light. The electrons themselves move at the speed of an ant. The whole of this chapter is about that crawl: how fast it is, what slows it down, how it heats the wire, how we steer it through circuits, and what happens when there are no electrons already in place and we must make our own.*

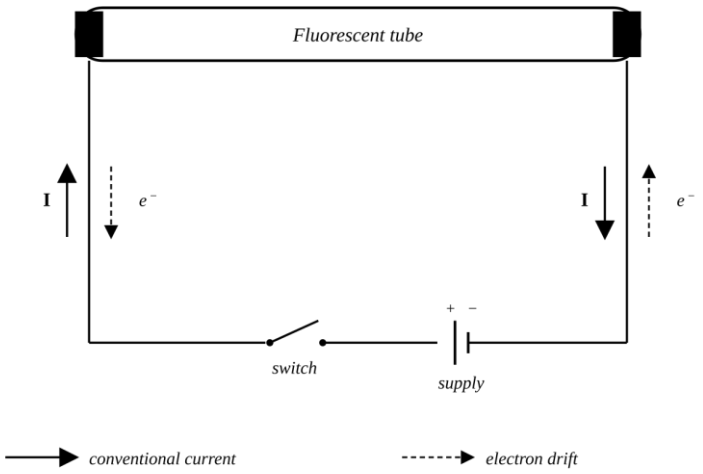


Figure: ***A fluorescent tube driven through a closed switch.*** *Solid arrows mark the conventional current I. Dashed arrows mark the electron drift, opposite to I on both wires.*

Chapter 11 ended with the time-varying current that flows briefly while a capacitor is charging or discharging. That current was the brief settling of charges into a new equilibrium; it lived for a few time constants and then died. The current of this chapter is different. It is what happens when the equilibrium is never reached at all, because a steady source keeps pushing charges through the circuit at the same rate forever. Once any initial transient has settled, every device in the world that runs on direct current is running on the steady current of this chapter, from the torch in a student's pocket to the immersion heater in a Nyamagana household to the power line that lit Mr. Akilikubwa's laboratory just now.

By the end of the chapter the reader will know what an electric current really is in microscopic terms, how Ohm's law arises from the collisions of drifting electrons with the lattice, what resistivity and conductivity and the temperature coefficient mean and where they come from, what an EMF is and why every battery has an internal resistance, how to combine resistors and cells in series and in parallel and in networks that bend under symmetry, how to convert a galvanometer into an ammeter or a voltmeter using only a shunt or a multiplier, how to apply Kirchhoff's two laws to any circuit that no series-and-parallel reduction can untangle, how the Wheatstone bridge and the potentiometer measure resistances and EMFs to a precision that no direct meter can match and why null methods are accurate at all, how electrical energy converts into heat and how to design a heating device that delivers exactly the power the design calls for, why TANESCO transmits power at hundreds of kilovolts

rather than at household voltage, what conditions a load resistor must satisfy to receive the maximum power a source can deliver, how a gas at low pressure can be made to conduct when its molecules are ionised, what distinguishes an initiated from a self-sustaining discharge, and how the light from a gas discharge tube carries a fingerprint of the gas inside.

The next twelve sections of this chapter answer those questions in the order they are best learned. We begin with the crawl itself.

ELECTRIC CURRENT AND THE DRIFT OF ELECTRONS

The chapter opener finished with electrons crawling at half a millimetre every second while the bulb lit immediately. This section puts the mathematics on that scene. To make a wire carry a steady current, three numbers have to cooperate: *how fast each electron drifts, how thickly the electrons fill the wire, and how big the cross-section of the wire is.* Once those three numbers are in hand, the current itself follows in a single line of algebra. The job of the section is to identify each of the three, name it, and put it in its place in that line.

What an electric current is

Imagine a copper wire with a steady current flowing through it. Pick any cross-section of the wire, a single slice perpendicular to its length. In one second a certain quantity of charge crosses that slice. In the next second the same quantity crosses again, and the second after that, and so on, as long as the source keeps pushing. The current I is the quantity of charge that crosses the slice per second:

$$\mathbf{I = \frac{Q}{t}}$$

where Q is the charge that crosses in time t. The unit of current follows from the definition. One ampere is one coulomb per second:

$$1\text{A} = 1\frac{\text{C}}{\text{s}}$$

If the current is not steady, the same idea applies to a vanishingly short interval of time rather than a finite one:

$$\text{I} = \frac{\text{dQ}}{\text{dt}}$$

We will use the steady form for almost all of this chapter. The varying form returns in Chapter 13 when we deliberately make the current change with time and watch the magnetic field follow it.

Charge does not come in arbitrary amounts. It comes in lumps of electronic charge $e = 1.6 \times 10^{-19}$ C. If N electrons cross our slice in time t, the total charge that has crossed is:

$$\text{Q} = \text{Ne}$$

and the current is therefore:

$$\mathbf{I = \frac{Ne}{t}}$$

A 1A current carries one coulomb per second past every cross-section of the wire. The number of electrons that have to cross to deliver that coulomb in one second is $N = 1\text{C} / (1.6 \times 10^{-19}\ \text{C}) \approx 6.25 \times 10^{18}$ electrons every second, more than six quintillion of them. That is the first surprise of the chapter. Even the smallest current of any practical interest moves more electrons through every cross-section every second than there are stars in every galaxy in the observable universe.

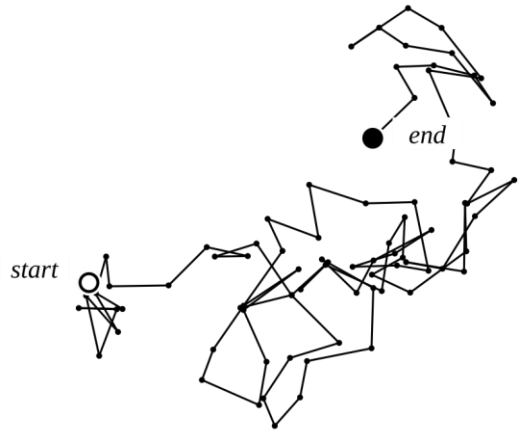
A note on direction:

Convention. Throughout this book, and on every figure that follows, the arrow drawn on a wire shows the direction of **conventional current**, which *is from the positive terminal of the source to the negative terminal through the external circuit.* Electrons in a metal drift in the **opposite** direction. We use the conventional direction in every formula. The electron-drift direction enters the discussion only when a question explicitly asks about it.

The convention is older than the discovery of the electron. Benjamin Franklin guessed in the 1750s that electric current was the flow of a single positive fluid from one place to another, and his choice stuck. We now know that in a metal the carriers are negative electrons. In a semiconductor the carriers may be electrons (negative) or holes (vacancies in the electron sea that behave as positive carriers and drift opposite to the electrons). In an electrolyte both positive and negative ions contribute, drifting in opposite directions. Regardless of what the carriers are, the conventional current measures the net flow of positive charge. The convention is therefore not wrong; it is simply about a different thing from the actual motion of the carriers.

The microscopic picture: random thermal speed and the small drift on top

Take a copper wire at room temperature with the switch open. The free electrons (about one per copper atom, the outermost electron that has parted with its atom) are not at rest. They fly about randomly at speeds of around 10^5 to 10^6 m/s, colliding with the lattice ions every few hundredths of a picosecond, ricocheting, colliding again, ricocheting again. Each electron is going somewhere very fast at every moment, but the swarm as a whole goes nowhere. The average velocity, taken over any single electron over any meaningful time, is exactly zero. The figure below sketches a typical electron's path through many collisions. The start and the end are close together; the motion in between is chaos.



no applied field: net displacement is small after many collisions

Figure: ***A single electron's path during many collisions with the lattice ions, in a metal with no applied electric field.*** *The motion is fast and chaotic, but the start and the end are close together: the net displacement is essentially zero.*

Now close the switch. An electric field E appears inside the wire, pointing along the wire's length. Each electron, in addition to its random thermal motion, picks up a tiny extra velocity in the direction opposite to E (because electrons are negative). This tiny added velocity is the **drift velocity**, written $\mathbf{v_d}$**.** The next figure shows what the same electron's path looks like now: the chaos is essentially unchanged, but a small systematic drift has been laid over it.

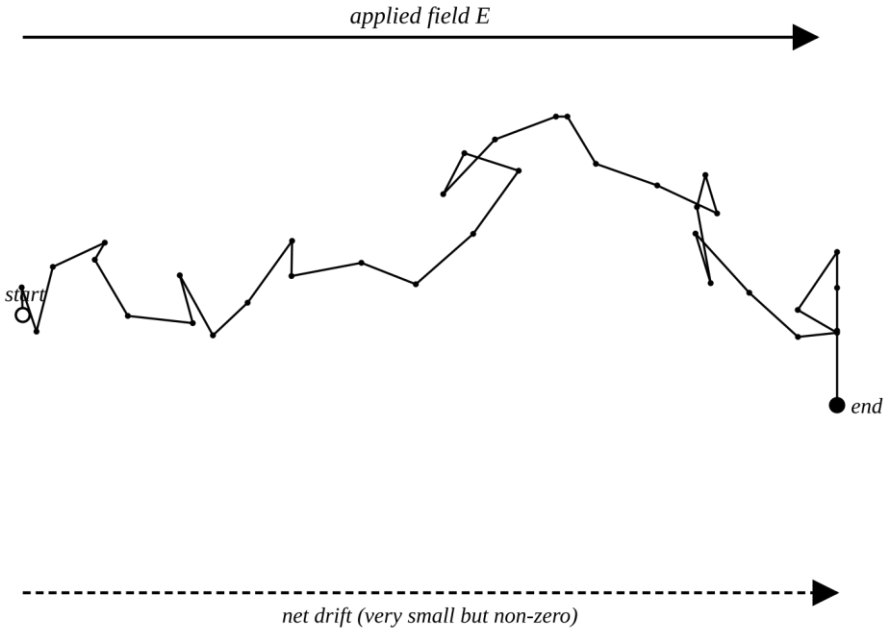


Figure: *With an applied field E along the wire, the same electron's chaos is now slightly biased.* ***After many collisions, the electron has drifted a small distance in one direction. The drift is the current.***

The drift velocity in a typical copper wire carrying a typical current is around 10^{-4} to 10^{-3} m/s, that is, a tenth of a millimetre to one millimetre every second. An ant carrying a leaf moves several centimetres every second. The drift is, almost literally, slower than the ant.

This is the central pedagogical point of the chapter opener and this section together. The drift is a tiny perturbation on top of a furious random motion. When the switch closes, the electric field fills the wire at the speed of light, and every electron everywhere in the wire begins drifting at the same instant. The bulb lights immediately because the wire is already full of electrons. No electron has to travel from the switch to the bulb; the field tells the electrons that already line the path between switch and bulb to start moving, and they do, **all at once.**

From the drift velocity to the current

We now have a picture of every free electron in the wire drifting uniformly with speed v_d in the conventional direction. To convert that picture into the current I, picture the wire as a cylinder of cross-sectional area A and length l, and call the number density of free electrons n, that is, the number of free electrons per cubic metre.

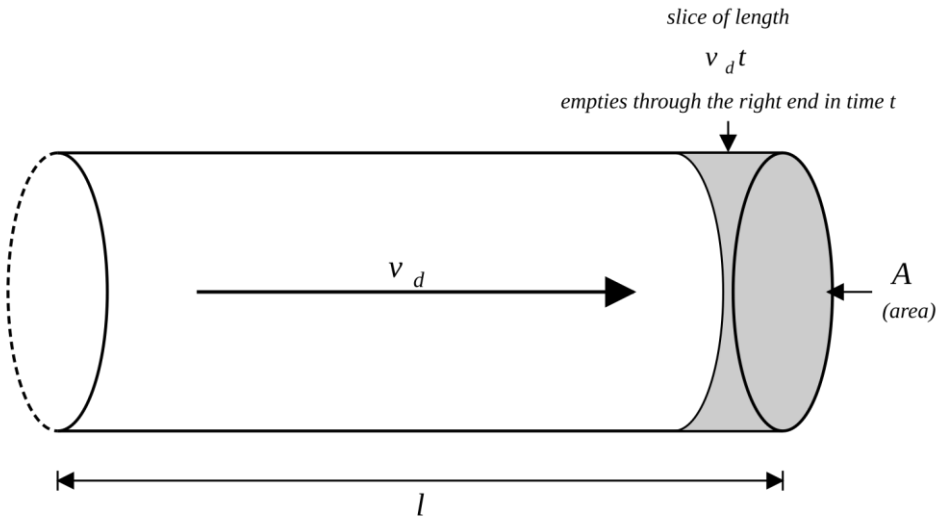


Figure: ***A cylindrical conductor used in the derivation.*** *Free electrons inside drift with speed v_d in the direction shown. The shaded slice at the right is the portion that empties through the right end during a time t.*

The total number of free electrons inside the cylinder is the number density multiplied by the volume of the cylinder:

$$N = nAl$$

In a time $t = \frac{l}{v_d}$ every electron in the cylinder has drifted out through the right end, assuming the drift is uniform throughout the cylinder, which to first order it is. The total charge that crosses the right end during that time is Ne. So the current is:

$$I = \frac{Ne}{t} = \frac{(nAl) \times e}{\left(\frac{l}{v_d}\right)} = neAv_d$$

The l in the numerator and the l in the denominator cancelled cleanly, and out fell a formula that does not depend on the length of the cylinder at all. It depends only on quantities that belong to the material (n and e) and to the cross-section of the wire (A) and to how fast the electrons are drifting (v_d). The result is the single most important formula of the section, the one every later calculation in this chapter starts from:

$$\mathbf{I = neAv_d}$$

Solving for the drift velocity:

$$\mathbf{v_d = \frac{I}{neA}}$$

Four ingredients sit on the right-hand side: how many free electrons per cubic metre the material has (n), the electronic charge (e), how big the cross-section is (A), and how much current you are driving (I). Three of the four belong to the wire itself; only one, the current, belongs to whatever is at the ends of the wire. Once we know what the wire is made of and how thick it is, the drift velocity is simply proportional to the current.

Current density

The combination I divided by A appears so often in what follows that it earns its own symbol and its own name. The **current density** is:

$$\mathbf{J = \frac{I}{A}}$$

with units of A/m². For our wire, substituting $I = neAv_d$:

$$\mathbf{J = nev_d}$$

Notice that J does not depend on the size of the wire. It depends only on what carriers are inside the material (n, e) and how fast they are drifting (v_d). *A thin wire and a thick wire made of the same material can carry different currents but the same J, provided their drift velocities are the **same**.* J is the intensive form of current; I is the extensive form. When in later sections we want to compare how easily a copper bar and a steel bar carry current, we will compare their values of J at the same applied field, never their values of I (which depend on the size of the bar).
A foreshadowing observation, to be earned properly in the next two sections: J is proportional to the electric field E that drives the current. In symbols:

$$\mathbf{J = \sigma E}$$

where σ is a property of the material called its conductivity. This is the local form of Ohm's law. It tells us that the current density at any point of any wire is set by the field at that point and by the conductivity of the material there, full stop. We will derive both of these from the microscopic collisions in the next section.

Drift versus thermal: the ratio that surprises everybody

We can now make the central comparison promised at the start of the section. The thermal speed of an electron in a metal at room temperature is about 10^5 to 10^6 m/s. The drift speed in the same metal carrying ordinary current is about 10^{-4} to 10^{-3} m/s. The ratio is:

$$\frac{v_{thermal}}{v_{drift}} \approx 10^8 \text{ to } 10^9$$

The drift is between a hundred million and a billion times slower than the random motion it is laid on top of. If you could film the electrons and slow the playback enormously, you would see chaotic flying in every direction; only after watching for a very long time, with a very steady eye, would you notice that on average the swarm has shifted by a tiny amount in one direction. That tiny average shift is the **current.**

And that completes the puzzle of the chapter opener. The bulb lights immediately because the wire was already full of electrons before the switch closed. When the switch closes, an electric field propagates through the wire at the speed of light, telling every electron at every point of the wire to start drifting. They all do, simultaneously. The bulb illuminates not because one electron rushed from the switch to it, but because the electrons that were already inside the bulb's filament started drifting at the same instant the electrons inside the switch did. The signal travels at the speed of light. The carriers crawl. The wire connects them because it was full all along.

Five examples follow. They each compute one of three things: how slow the electrons drift, how thickly the current packs through the cross-section, or how many electrons sneak past a given point each second. The numbers you are about to meet are absurd in both directions: miraculously slow and miraculously numerous. Once you have written them down, the queue of buckets from the chapter opener will have earned its keep.

BINDER Example 1

A current of 5A flows through a copper wire whose cross-sectional area is 2mm². The number density of free electrons in the copper is taken as 10^{28} m⁻³. Find the drift velocity of the free electrons in the wire.

Solution
The drift velocity is given by:

$$v_d = \frac{I}{neA}$$

Substituting:

$$v_d = \frac{5A}{(10^{28}\ m^{-3})(1.6 \times 10^{-19}C)(2 \times 10^{-6}\ m^2)} = 1.5625 \times 10^{-3}\ m/s$$

Making Sense of the Answer: *A drift velocity of 1.56×10^{-3} m/s is roughly 1.5 millimetres every second. If you asked a single electron to walk from a wall socket to a lamp three metres away, at this drift speed the trip would take about 1900 seconds, more than half an hour. The lamp does not wait half an hour to light because the electrons inside its filament were already there when the switch was closed. The lamp lights at the speed of light.*

Think Like a Physicist: *Whenever a question asks for a drift velocity, expect an answer in the 10^{-4} to 10^{-3} m/s range for ordinary currents and ordinary wires. If you compute a drift velocity larger than 1m/s, somewhere in your arithmetic you have lost a factor of about 10^3, and you should hunt the missing thousand down before submitting the answer. Sanity always before substitution.*

BINDER Example 2

A copper wire of cross-sectional area 4mm² and length 4m carries a current of 10A. The number density of free electrons in the copper is 8×10^{28} m⁻³. Find: (a) the drift velocity of the free electrons in the wire, (b) the time required for an individual electron to drift the entire length of the wire.

Solution
(a) From the same drift-velocity formula:

$$v_d = \frac{I}{neA}$$

Substituting:

$$v_d = \frac{10A}{(8 \times 10^{28}\,m^{-3})(1.6 \times 10^{-19}C)(4 \times 10^{-6}\,m^2)} = 1.95 \times 10^{-4}\,m/s$$

(b) The transit time of an individual electron is the length of the wire divided by its drift velocity:

$$t = \frac{l}{v_d}$$

$$t = \frac{4m}{1.95 \times 10^{-4}\,m/s} = 2.05 \times 10^{4}s \approx 5.70\ hours$$

Making Sense of the Answer: *An individual electron takes nearly six hours to drift the four metres from one end of the wire to the other. Meanwhile the electric signal that drives the current propagates through that same four-metre wire in roughly 4m / (3 × 10⁸ m/s) ≈ 1.3 × 10⁻⁸ s, about thirteen nanoseconds. The signal beats the electron by a factor of roughly (5.70 × 3600) / (1.3 × 10⁻⁸) ≈ 1.6 × 10¹², about twelve orders of magnitude. The signal flies. The carriers crawl. The wire connects them because it was full all along.*

Think Like a Physicist: *Transit-time questions about drifting electrons are designed to surprise. The intuition for an ordinary copper wire is that an individual electron makes the trip in hours, while the electric signal that drives the current makes the trip in nanoseconds. If your answer for the transit time of an individual electron in a sensibly thick copper wire comes out in milliseconds, you have probably confused signal speed with drift speed; redo the calculation.*

BINDER Example 3

A copper wire has cross-sectional area 1.5mm² and carries a current of 6.0A. Taking the number density of free electrons in copper as $8.5 \times 10^{28}\ m^{-3}$, find:

(a) the current density J in the wire,

(b) the drift velocity of the free electrons.

Solution

(a) From the definition J = I/A:

$$J = \frac{I}{A} = \frac{6.0A}{1.5 \times 10^{-6}\,m^2} = 4.0 \times 10^{6} A/m^2$$

(b) From:

$$v_d = \frac{J}{ne}$$

Substituting:

$$v_d = \frac{4.0 \times 10^{6} A/m^2}{(8.5 \times 10^{28}\,m^{-3})(1.6 \times 10^{-19}C)} = 2.94 \times 10^{-4}\,m/s$$

Making Sense of the Answer: *A current density of 4 × 10⁶ A/m² sounds enormous until you remember that the cross-section is only 1.5mm², so the actual current of 6A is what a household lamp draws. The same J in a wire of cross-section 10mm² would still be 4 × 10⁶ A/m² because J does not care how thick the wire is. The drift velocity is again in the familiar 10⁻⁴ m/s range, exactly the order of magnitude the section promised.*

Think Like a Physicist: *Whenever a question hands you a current and a cross-section together, compute J first, then v subscript d if you need it. J is the intensive quantity; it cleans up the algebra by removing the cross-sectional area from your formulas. After J, the drift velocity is one division away.*

REAL Example 4

It is past ten o'clock at night in Miono. Kipute has been reading her notes for two hours under the 100W ceiling lamp in her study room. She reaches up to adjust the angle of the shade, and her hand brushes the supply wire near the lamp holder. The wire is warm. Not hot, not dangerous, but warm enough to notice.

Kipute: *Sir, the wire to my reading lamp felt warm after two hours under it last night. Does that mean the electrons inside it are moving fast?*

Mr. Akilikubwa: *Let us answer that with a number rather than an opinion. Your lamp draws its 100W of power from the 240V mains; the supply wire is ordinary household copper with cross-sectional area 1.5mm² and free-electron number density 8.5 × 10²⁸ m⁻³.* ***Find the drift velocity of the free electrons in that wire, and then we shall talk about whether they are moving fast.***

Help Kipute:

(a) to calculate the drift velocity of the free electrons in the supply wire to her reading lamp, and
(b) from the answer in (a), explain to her why a wire that feels warm to the touch is not a wire full of fast electrons.

Solution

(a) The drift-velocity formula needs the current I, but the question hands us a power and a voltage instead. Power equals voltage times current, so:

$$I = \frac{P}{V}$$

Substituting to $v_d = \frac{I}{neA}$:

:

$$v_d = \frac{P}{VneA} = \frac{100W}{240V \times (8.5 \times 10^{28}\ m^{-3}) \times (1.6 \times 10^{-19}C) \times (1.5 \times 10^{-6}\ m^2)} = 2.04 \times 10^{-5}\ m/s$$

That is about twenty micrometres every second, twenty thousandths of a millimetre every second.

Have you noticed this? Twenty micrometres every second is slower than a snail crawling on tar at noon. An individual electron in the wire to Kipute's reading lamp would need nearly an hour to drift the width of her fingertip. The electrons are not moving fast in any sense that the word ordinarily carries. And yet the wire is warm to the touch after two hours of operation. The warmth, therefore, cannot be a consequence of the electrons moving fast, because the electrons are not moving fast. It must come from somewhere else.

That somewhere else is the collisions. Now, let us answer part (b).

(b) The warming arises from collisions between the electrons and the lattice ions. Every time a drifting electron meets a lattice ion, it surrenders a tiny piece of the velocity it had just picked up from the field, and the lattice absorbs that velocity as a small vibration of the ion. With more than $10^{28}\ m^{-3}$ free electrons in every cubic metre of copper, slamming into the lattice ions many times every picosecond, those small vibrations add up to a wire that is detectably warm after long operation.

The mechanism is the collision-induced heating that we call Joule heating, and we will devote a whole section to it later in this chapter. For now Kipute can trade her worry for a corrected intuition: a warm wire is not a wire full of fast electrons. **It is a wire full of slow electrons colliding furiously.**

Making Sense of the Answer: *Both halves of the question concluded with the same moral: the carrier speeds inside an ordinary household wire are absurdly slow, and the macroscopic things the wire seems to be doing (lighting a room instantly when the switch closes, warming after long use) have nothing to do with electrons moving fast. The chapter opener settled the lighting; this example settles the warming. In each case the answer is the same: the wire was already full of carriers; the electric field set them all drifting simultaneously; their slow drift gets dissipated into heat by collisions with the lattice. Slow electrons, fast field, immediate response, slow heating. Four pieces, one story.*

Think Like a Physicist: *Whenever a question hands you a power rating and a supply voltage instead of a current, the first equation to reach for is I = P/V. From there, the drift velocity is exactly one substitution away. Power-voltage data hides the current; the current hides the drift; and the drift is usually what the question is really after. Three steps, not one.*

HOT Example 5

A steady current passes through a cross-section of a conducting wire. Over a period of 30 seconds, 7.5×10^{19} electrons cross the cross-section. Find:

(a) the total charge that has crossed the cross-section in the 30s,
(b) the magnitude of the current in the wire,
(c) explain briefly why a lamp connected in series with this wire would light up the instant the switch is closed, even though the individual electrons in the wire travel at drift velocities of only millimetres per second.

Solution

(a) The total charge is given by:

$$Q = Ne = (7.5 \times 10^{19}) \times (1.6 \times 10^{-19}C) = 12C$$

(b) The current is the charge that crossed per unit time:

$$I = \frac{Q}{t} = \frac{12C}{30s} = 0.4A$$

(c) The electrons that are about to enter the lamp's filament were already there before the switch closed; they did not have to travel from the source to the filament. When the switch closes, the electric field fills the wire at the speed of light, telling every electron, including those already inside the filament, to start drifting at the same instant, so the filament heats and lights immediately because the electrons inside it began drifting immediately, not because an electron rushed across the room to reach it.

Making Sense of the Answer: *Twelve coulombs in thirty seconds gives 0.4A, roughly the steady current drawn by a small household lamp. The pedagogical part (c) is the real point of the question. The lamp's response is set by the speed of the field, not by the speed of the carriers, and that point separates students who understand current from students who only know how to compute it.*

Think Like a Physicist: *In any question that gives you a number of electrons over a time interval, the conversion is always Q = Ne and I = Q/t. The two facts merge into one statement: I times t gives the count N = It/e of electrons that have crossed. Many questions in this chapter can be unlocked by spotting which of these three quantities (Q, I, or N) is the implicit one in the question.*

HOT Example 6

A current of 0.5A flows through a 40W lamp. Calculate the number of electrons that flow through the filament of the lamp in one second.

Solution

The number of electrons that carried this charge is:

$$N = \frac{Q}{e} = \frac{It}{e}$$

Substituting:

$$N = \frac{(0.5A)(1s)}{1.6 \times 10^{-19}C} = 3.125 \times 10^{18} \text{ electrons}$$

Making Sense of the Answer: *This is more than three quintillion of electrons. The filament has three quintillion electrons entering and three quintillion leaving every second, and the difference between entering and leaving is exactly zero by charge conservation. The current is not the electrons going in (or out); it is the rate at which they cross.*

Think Like a Physicist: *When a question gives you both a current and a power rating, ask which of the two the question actually needs. Examiners regularly hand students decorative data to test whether they recognise the path to the answer. If the question asks for electrons per second, the answer depends only on I and t; the wattage is bait.*

Six examples done. The numbers should now sit comfortably in mind. Drift velocity in tenths of a millimetre every second. Current density in millions of amperes every square metre. Electron count past any cross-section in quintillions every second. The crawl is mathematically respectable; if anything, it is too slow to be respectable. The next question is what makes one wire crawl easier than another, why a copper wire does the job better than a steel one, and why a hot wire does it slightly worse than a cold one. That is Ohm's law, and the next section walks straight to it.

OHM'S LAW FROM MICROSCOPIC THEORY

The last section answered how a current I flows through a wire once the carriers are drifting. It left one question hanging, the one any honest student must ask next: *what makes the carriers drift in the first place, and why does one wire pass more current than another?* Two answers are on the table. The first is the answer Georg Simon Ohm wrung out of his apparatus in 1827, after wrapping platinum wires of different lengths around wooden frames and passing currents through them from a thermopile. It is empirical, brief, and astonishing in its simplicity. The second is the answer the microscopic picture of the previous section quietly built for us without our noticing. The two answers agree, and the agreement is the content of this section.

Ohm's law as an empirical regularity

Take a copper wire of fixed length and thickness. Connect it across a variable voltage source. Keep the wire in a beaker of cool water so its temperature does not change as current passes through it. Slowly raise the voltage from zero, and at each setting measure the current the wire carries. Plot the current I against the voltage V. The result is a straight line that passes through the origin and extends with the same slope into the negative-voltage quadrant if the source is reversed.

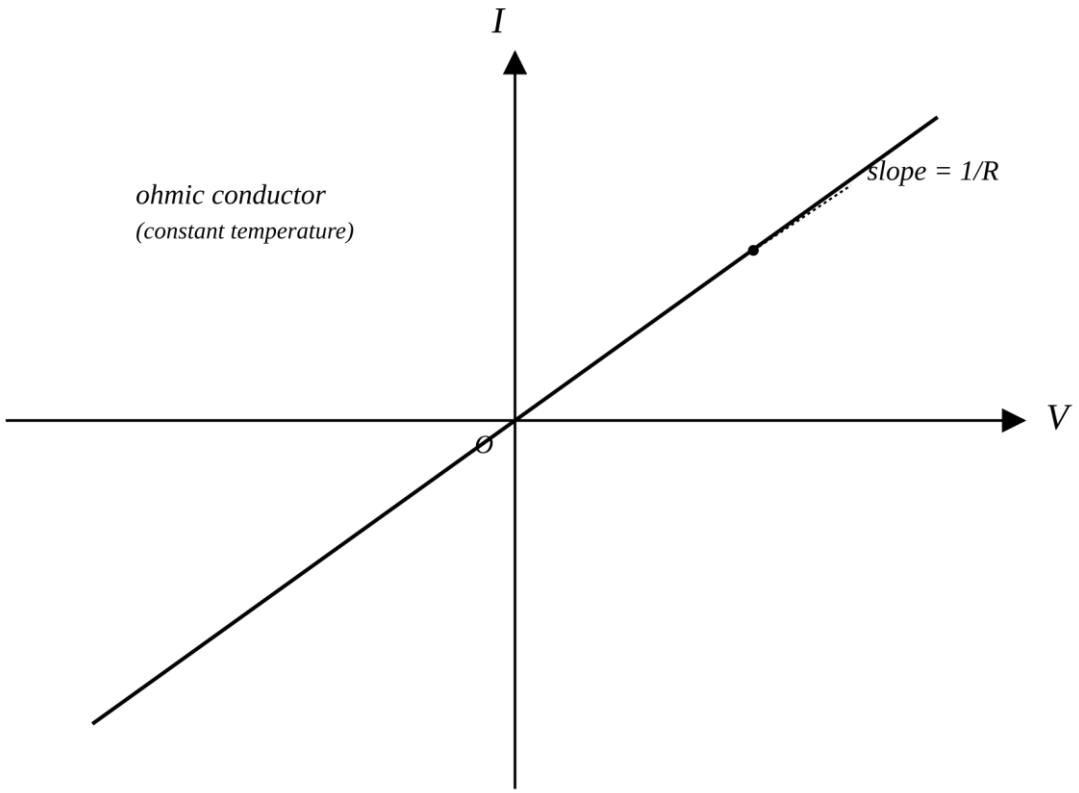


Figure: ***Current–voltage characteristic of an ohmic conductor at constant temperature.*** *The graph is a straight line through the origin in both polarities, and the reciprocal of its slope is the resistance of the conductor.*

Ohm's law states that:

The current that flows through a conductor between two points is directly proportional to the potential difference applied across the two points, provided physical conditions, especially the temperature, remain constant.

Mathematically:

$$V \propto I$$

Calling the constant of proportionality R, Ohm's law becomes:

$$V = IR \text{ or } \mathbf{I} = \frac{\mathbf{V}}{\mathbf{R}}$$

The constant **R** is called the **resistance** of the conductor, and Ohm's law itself defines it: the **resistance** of a piece of wire *is the voltage across the wire divided by the current through it, when the wire is held at constant temperature*. In symbols:

$$R = \frac{V}{I}$$

The unit of resistance is the volt per ampere, given the special name ohm with the symbol Ω:

$$1\Omega = 1V/A$$

A wire of resistance 1Ω carries one ampere of current under a voltage of one volt across its ends. A wire of resistance 10Ω carries only one tenth as much current under the same voltage. A wire of resistance 0.1Ω carries ten times as much. The resistance is the wire's individual contribution to the question of how much current the voltage pushes through it.

Two important warnings about the law itself before we proceed.

- The law is not a law of nature in the sense that Newton's laws of motion are. It is an empirical regularity that happens to be true for some materials and false for others. *Materials that obey it* are called **ohmic**. Most metals at fixed temperature are ohmic. Copper sulphate solution is ohmic. Carbon at fixed temperature is approximately ohmic. *The conductors that disobey it* are called **non-ohmic.** The filament of an incandescent lamp at varying current is non-ohmic, because the filament heats as the current rises, and its resistance changes with temperature. The junction of a semiconductor diode is strongly non-ohmic, because it lets current through easily in one polarity and almost not at all in the other.

We will sketch the non-ohmic curves at the end of the next section, after we have understood why temperature changes the resistance. For now, every conductor in the worked examples that follow is ohmic, and Ohm's law holds in the simple form just given.

- Ohm announced his result as an empirical regularity, true in the laboratory but lacking any deeper explanation. For the better part of seventy years the physics community took the law at face value, valid where the experiments said it was valid and silent on the question of why. The explanation arrived only at the very end of the nineteenth century, with the discovery of the electron and the construction of the microscopic model we developed in the previous section.

We will now follow that explanation through, in four steps. At the end of the four steps Ohm's law will reappear, not as an empirical truth but as a consequence of more elementary physics.

The microscopic question

From the previous section we have an exact relation between the current in a wire and the drift velocity of its free electrons:

$$I = neAv_d$$

Three of the four factors on the right are properties of the wire that do not depend on whether a voltage has been applied: the number density of free electrons n, the electronic charge e, and the cross-sectional area A. The fourth factor, the drift velocity v_d, does depend on the applied voltage; it is in fact the only place the voltage can enter the expression. To understand how I depends on V, we therefore need to understand how the drift velocity depends on the applied electric field. The next two steps build that relationship.

An electron between collisions

Imagine a single free electron deep inside the wire, mid-flight between two collisions with the lattice. The applied electric field E points along the wire. The force on the electron is the field times the electron's charge, but because the electron is negative, the force points opposite to the field. In symbols (taking e for the magnitude of the electronic charge, so $e > 0$):

$$F = eE$$

where the equality is taken in magnitude, and the force points opposite to the field.

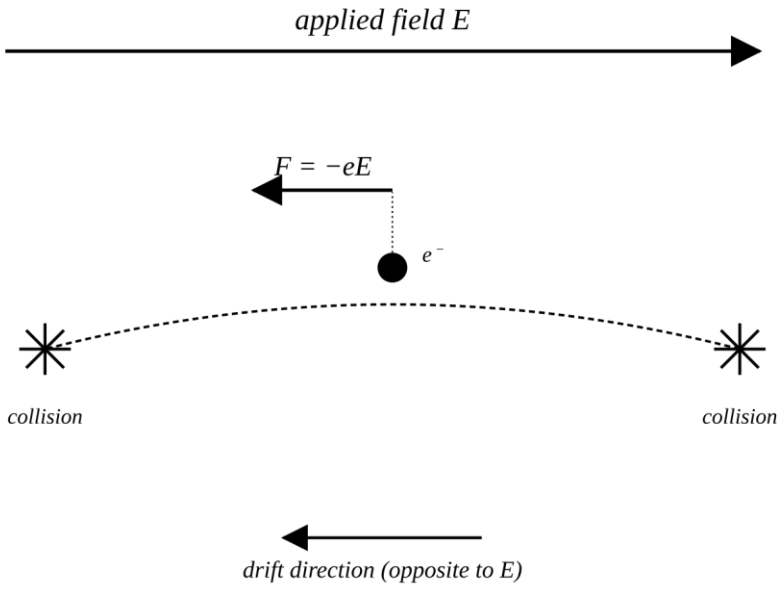


Figure: ***One free electron during a single flight between two collisions with the lattice***. *The applied field points to the right, the force on the electron points to the left, and the electron picks up a small drift velocity to the left during the free flight before the next collision randomises its velocity.*

Step one: the acceleration. By Newton's second law, the acceleration of the electron during the free flight is the force divided by the electron's mass m:

$$a = \frac{F}{m} = \frac{eE}{m}$$

The acceleration points opposite to the field. The negative sign that this would carry in a vector treatment has been absorbed into our convention of always working with magnitudes when discussing drift; the direction is built into the conventional current picture and is not in dispute.

Step two: the relaxation time. If the electron were free to accelerate forever, it would reach speeds approaching the speed of light in a fraction of a second under any electric field of laboratory strength. It does not, because every so often it smashes into a lattice ion (or an impurity, or a vibrating lattice defect) and loses all memory of the drift it had just acquired. The average

time between two such collisions is called the **relaxation time** of the metal, written **τ**. Starting from **zero** drift just after one collision and accelerating for an average time τ before the next, the electron acquires an average drift velocity:

$$v_d = a\tau = \frac{eE\tau}{m}$$

The drift velocity is the acceleration multiplied by the average time the acceleration is allowed to operate. **Putting it together** ***Step three: substitute into $I = neAv_d$.*** Replacing the drift velocity in the previous section's expression by the microscopic formula just derived:

$$I = neA \times \frac{eE\tau}{m} = \frac{ne^2A\tau}{m} \times E$$

For a uniform conductor of length l carrying a steady current, the field inside it is the voltage across it divided by its length:

$$E = \frac{V}{l}$$

Substituting once more:

$$\mathbf{I = neA \times \frac{eE\tau}{m} = \left(\frac{ne^2A\tau}{ml}\right) \times V}$$

Step four: read off the resistance. The current is now proportional to the voltage. The constant of proportionality is a single bracketed expression involving only properties of the material (n, e, m, τ) and of the geometry (l, A). Comparing with the empirical Ohm's law $I = \frac{V}{R}$, the resistance must be the reciprocal of that bracketed expression:

$$\mathbf{R = \frac{ml}{ne^2A\tau}}$$

This is the microscopic expression for the resistance of a wire. Notice that it has been built from quantities every one of which we have already met: the electron's mass, the length and cross-sectional area of the wire, the number density of free electrons in the material, the electronic charge, and the relaxation time. Nothing in the formula refers to V or I; *the resistance is a property of the wire by itself,* ***not*** *of the experiment we choose to do with it.* That is exactly what Ohm's law required.

Reading the microscopic formula

The microscopic resistance formula separates cleanly into two parts.

On the one hand it depends on the **geometry** of the wire through the length and the cross-sectional area. **A longer wire has more resistance**, in direct proportion to its length, *because the electron must traverse more lattice before reaching the other end.* **A thicker wire has less resistance**, in inverse proportion to its area, *because the same current can be carried by more electrons* each drifting more slowly.

On the other hand the formula depends on the **material** of the wire through four atomic constants: the free-electron number density, the electronic charge, the electron's mass, and the relaxation time. The first three are the same for any metal that we happen to consider, more or less; copper, silver, aluminium, and iron all have free-electron densities of order 10^{28} to 10^{29} per cubic metre. The fourth, the relaxation time, varies wildly from material to material and is the property that distinguishes good conductors from poor ones.

It is conventional to separate the material part of the resistance from the geometric part by defining the **resistivity ρ**:

$$\mathbf{\rho = \frac{m}{ne^2\tau}}$$

in terms of which the resistance of a wire is:

$$\mathbf{R = \frac{\rho l}{A}}$$

The resistivity has units of Ωm and is a *property of the material alone*, independent of how the material is shaped into a wire. The resistance is the geometry multiplying the resistivity. The resistivity is the material; the resistance is the material wrapped in a particular shape. We will return to the resistivity in the next section, where we tabulate its values for the materials that appear in this book and study how it varies with temperature.

One more reading. We can now say precisely what makes a conductor ohmic. **A conductor is ohmic** *when τ depends on the temperature of the lattice but not on the field or the voltage applied across the conductor.* When that is so, the right-hand side of the microscopic formula is independent of V, and R is a constant; the current is then strictly proportional to the voltage. A conductor is non-ohmic when τ depends on the field as well, either because the field is so strong that it begins to heat the conductor (as in an incandescent lamp) or because the conductor's structure responds nonlinearly to the field (as in a semiconductor diode). Both effects produce a voltage-dependent resistance, and the I-V curve curves away from a straight line. The microscopic picture explains in one stroke why the ohmic class is so wide (most metals at fixed temperature, where τ is set by the lattice and not by the field) and why the non-ohmic class is so important (every device that does something interesting with current, from a transistor to a filament lamp).

Conductance

Conductance *is the reciprocal of resistance:*

$$G = \frac{1}{R}$$

Its unit is the siemens (S), defined as the reciprocal of the ohm: 1S = 1/Ω = 1A/V. Where **resistance answers the question** *"how much voltage do you need to push a given current through me*", **conductance answers** *"how much current do you get out of me for a given voltage".* Both descriptions are equivalent; the choice is stylistic. Most of this book uses R, but conductance occasionally simplifies the algebra in parallel networks, where the total conductance is the sum of the individual conductances of the branches. We will meet that simplification later in the chapter, when we come to circuit networks.

Ohm's law is not Joule's law

A common student error is to merge Ohm's law with Joule's law into a single rule. They are different laws about different things. Ohm's law $I = \frac{V}{R}$ tells you how the current through a conductor relates to the voltage across it. **Joule's law:**

$$\mathbf{P = I^2R}$$

tells you *the rate at which the conductor dissipates electrical energy as heat*. Ohm's law is empirical and approximate; **Joule's law** is **exact** for any conductor, ohmic or not, *because it follows directly from the definition of work done against a potential difference*. A filament lamp is non-ohmic; Ohm's law fails for it; $P = I^2R$ still holds. The full treatment of Joule heating waits until later in the chapter; for now the only point is that the two laws are not the same.

Three examples follow.

BINDER Example 7

A 12V battery is connected across a resistor of resistance 4Ω. Find the current that flows through the resistor.

Solution

Direct application of Ohm's law:

$$I = \frac{V}{R} = \frac{12V}{4\Omega} = 3A$$

Making Sense of the Answer: *Three amperes from a 12V battery across a 4Ω resistor. A current of this size is what a small car interior light draws, a couple of orders of magnitude smaller than the starter motor (which pulls hundreds of amperes for a brief moment) and a couple of orders of magnitude larger than a transistor radio. Three amperes is also what a dry cell can comfortably deliver for a few minutes before its internal resistance starts to matter.*

Think Like a Physicist: *Whenever a question gives you a voltage and a resistance and asks for a current, the answer is one line of arithmetic. The trick is not to overthink it. Ohm's law is the shortest equation in the chapter and the most reliable.*

REAL Example 8

It is half past five in the evening in Morogoro. Kipanga's mother has just put the kettle on for tea, and within a minute the lamp on the ceiling dims to about half its usual brightness. A few seconds later the kettle's indicator light fades. Kipanga, who has been working through a physics holiday package at the kitchen table, looks up. The mains voltage has dropped; this is the third brownout this week.

Kipanga: *Mother, the kettle will still boil eventually, will it not? The voltage has only dropped, it has not been cut off entirely.*

Mother (half-smiling)**:** *Eventually, yes. But it will take longer. Ask the kettle why.*

Help Kipanga:

(a) to calculate the resistance of the kettle's heating element, given that it is rated 1500W and is designed to run on the 240V mains;
(b) to predict the current the kettle now draws when the brownout has dropped the mains voltage to 200V, treating the heating element as ohmic at constant temperature;
(c) to calculate the power the kettle now dissipates in the brownout, and explain to Kipanga, why his mother predicted that the water will boil but more slowly.

Solution

(a) From the rated power and voltage, the rated current is:

$$I = \frac{P}{V}$$

Applying Ohm's law in the form R = V/I:

$$R = \frac{V}{I} = \frac{V}{\frac{P}{V}} = \frac{V^2}{P} = \frac{(240V)^2}{1500W} = 38.4\Omega$$

(b) Treating the heating element as ohmic at constant temperature, the resistance does not change when the voltage changes. The new current is the new voltage divided by the same resistance:

$$I = \frac{V}{R} = \frac{200V}{38.4\Omega} = 5.21A$$

(c) The power dissipated by the heating element at the lower voltage is:

$$P = VI = 200V \times 5.21A = 1042W$$

This is about seventy percent of the kettle's rated 1500W output.

The kettle will still boil the water because it still supplies sufficient power to raise the temperature of the two litres to boiling point, but it does so more slowly since the time taken to boil is inversely proportional to the power; therefore, with reduced power, a longer time is required.

Making Sense of the Answer: *Drop the voltage to five sixths of its rated value and the current drops to five sixths of its rated value (because the kettle is ohmic), but the power drops to twenty-five over thirty-six, about seventy percent of its rated value. The power drops faster than the voltage because power is voltage multiplied by current, and both factors are dropping in proportion to the voltage. This is the general fact: ohmic conductors dissipate power as the square of the voltage applied across them.*

Think Like a Physicist: *Whenever a kettle, heater, geyser, immersion rod, or iron is fed a voltage different from its rating, expect three quantities to change in lockstep: current scales as the voltage, power scales as the square of the voltage, and time to boil or to heat scales inversely as the square of the voltage. The fastest sanity check on any such question is to ask what fraction of the rated voltage is being supplied; everything else follows by squaring or square-rooting that fraction.*

HOT Example 9

The free electrons in copper have an average relaxation time 2.5×10^{-14}s at room temperature. Taking the number density of free electrons as $8.5 \times 10^{28}\ m^{-3}$, the electron's charge as $e = 1.6 \times 10^{-19}$C, and the electron's mass as $m = 9.1 \times 10^{-31}$kg, calculate the resistivity of copper predicted by the microscopic model and compare it with the tabulated value of $1.72 \times 10^{-8}\ \Omega$m.

Solution

The microscopic formula for the resistivity is:

$$\rho = \frac{m}{ne^2\tau}$$

Substituting:

$$\rho = \frac{9.1 \times 10^{-31}\ kg}{(8.5 \times 10^{28}\ m^{-3}) \times (1.6 \times 10^{-19}C)^2 \times (2.5 \times 10^{-14}s)} = 1.67 \times 10^{-8}\ \Omega m$$

Compare with the tabulated value of $1.72 \times 10^{-8}\ \Omega$m. The prediction is extremely close, differing by only about:

$$\left(\frac{1.72 \times 10^{-8}\ \Omega m - 1.67 \times 10^{-8}\ \Omega m}{1.72 \times 10^{-8}\ \Omega m}\right) \times 100\% = 2.9\%,$$

which shows that the microscopic model gives an excellent description of electrical conduction in copper.

Making Sense of the Answer: *About three percent agreement between a calculation built from four atomic-scale constants and a measurement taken in a copper wire at room temperature, with no fitting and no fudge factors, is one of the small triumphs of nineteenth-century physics. The free-electron model has its faults, particularly when temperatures get high or fields strong, but at room temperature in ordinary copper it is essentially right.*

Think Like a Physicist: *Whenever a derivation produces a number you can compare with a measurement, do the comparison. Numerical sanity checks are the fastest way to learn whether the microscopic model you have just spent a section building is worth believing. If a derivation predicts a quantity to within an order of magnitude of its measured value, the underlying physics is approximately right. If it predicts the quantity to within a few percent, the underlying physics is almost certainly right.*

Three examples done. Ohm's law has been stated, derived, applied to a battery and a resistor, applied to a kettle and a brownout, and finally checked against the resistivity of copper at the atomic scale. The microscopic formula has earned the chapter's trust. What remains is to follow it where it leads: *how does the resistivity differ between materials, why does copper conduct better than iron, and why does any metal's resistance rise as it heats up?* Those are the questions of the next section, and the answer in every case lives in the relaxation time.

RESISTIVITY, CONDUCTIVITY, AND TEMPERATURE DEPENDENCE OF RESISTANCE

The closing paragraph of the last section split the resistance of a wire into two halves and pointed at the second one. The first half belongs to the wire's geometry. The length and the cross-sectional area can be measured with a ruler and a calliper, and they are exactly the same numbers whether the wire is copper, steel, nichrome, or glass. The second half belongs to the wire's material. The number density of free electrons, the electronic charge, the electron's mass, and the relaxation time are properties of the substance the wire is made of, and no ruler will tell you their values. Copper conducts better than iron because of this second half. Silver conducts even better than copper, aluminium only slightly worse, and fused quartz very much worse. None of those facts depend on how thick or thin you have drawn the wire. They depend on what you have drawn the wire out of.

The cleanest way to handle this split is to write the material half down once, give it a name, and then forget the geometry until you need it again. The material half is called the resistivity of the substance, and the next several pages are about how to read it, where to find it, and how it changes when the wire is heated.

The resistivity of a substance

Imagine, first, a very fat and very short piece of the substance: a cube one metre on a side, with current entering through one face and leaving through the opposite face. The length of this block is one metre. Its cross-sectional area is one square metre. If you measured the resistance between the two faces, you would be measuring a quantity that depends only on what the block is made of. Lengthen the block to two metres, with the same cross-section, and the resistance doubles. Halve the cross-section, with the length still one metre, and the resistance doubles again. So the resistance of a block of given material is proportional to the length and inversely proportional to the cross-section. What the proportionality is, in ohms per metre divided by square metres, is a number that is set by the substance itself. That number is the **resistivity**, and we write it with the Greek letter **ρ**. It can be defined as *a property of a material that measures how strongly the material opposes the flow of electric current*

From the microscopic derivation of the last section, the resistance of a wire of length l and cross-sectional area A was:

$$R = \frac{ml}{ne^2A\tau}$$

Four of the six quantities on the right belong to the material (the electron's mass m, the number density of free electrons n, the electronic charge e, and the relaxation time τ); the other two belong to the geometry (the length l and the cross-sectional area A). Group the material quantities together by writing:

$$\rho = \frac{m}{ne^2\tau}$$

Then the resistance becomes simply:

$$\mathbf{R = \frac{\rho l}{A}}$$

The new quantity ρ is called the resistivity of the substance. Its unit follows directly from the formula by transposing for ρ:

$$\rho = \frac{RA}{l}, \qquad \text{unit: } \frac{\Omega m^2}{m} = \Omega m$$

The resistivity is therefore measured in ohm metres. A substance with a resistivity of $1\Omega m$ would offer a resistance of one ohm to a block one metre long and one square metre in cross-section. In practice no ordinary conductor comes close to that; copper, the workhorse of household wiring, has a resistivity of about $1.7 \times 10^{-8}\Omega m$ (i.e. 1.7 nanoohm metres), ten billion times smaller. The unit is large because most useful conductors are very good at conducting.

Two things can now be said precisely. The resistance is what the wire offers; the resistivity is what the substance offers. The resistance depends on how the substance has been shaped into a wire; the resistivity does not. Two copper wires of different thicknesses have different resistances but the same resistivity. If you needed a single number to describe how good a conductor a substance is, independent of the shape it has been drawn into, the resistivity is that number.

Conductivity

When we want to say a substance is a good conductor, we often prefer a number that grows when the substance conducts well, rather than one that shrinks. The reciprocal of the resistivity does the job, and it is called the **conductivity**, written with the Greek letter **σ**:

$$\boldsymbol{\sigma = \frac{1}{\rho}}$$

Substituting the microscopic expression for ρ:

$$\boldsymbol{\sigma = \frac{ne^2\tau}{m}}$$

Conductivity *is a property of a material that measures how easily electric current flows through it.*

The unit of conductivity is the reciprocal of the unit of resistivity, which is the reciprocal ohm metre, more often called the siemens per metre and written S/m. A conductivity of one siemens per metre is rather poor; copper reaches roughly $6 \times 10^7 S/m$ (sixty million siemens per metre), which is the standard against which every other conductor is measured.

The conductivity gives Ohm's law its tidiest form. From the previous section we already had $J = \sigma E$ foreshadowed; we can derive it now in one line.

Starting from the resistance form $V = IR$ and substituting $V = El$ (for a uniform electric field along a wire of length l) and $R = \frac{\rho l}{A}$:

$$El = I\frac{\rho l}{A}, \quad \text{hence} \quad \frac{I}{A} = \frac{E}{\rho}, \quad \text{that is,}$$

$$\mathbf{J = \sigma E}$$

The current density at any point in any conductor equals the conductivity at that point multiplied by the electric field at that point. This is the local form of Ohm's law. It removes the geometry of the wire from the discussion altogether and is the form physicists prefer when working with materials whose cross-sections vary along the length, or whose fields vary in space. The integral form $V = IR$ is the same statement integrated over the length of a uniform conductor. Both versions of Ohm's law follow from the same microscopic picture; the difference is only which scale you choose to look at.

Reading the resistivity of common materials

The next figure tabulates the resistivity, at room temperature, of eleven materials that turn up in this book and in the engineering of any Tanzanian household. The axis is logarithmic because the values stretch across twenty-five orders of magnitude, from ten-billionths of an ohm metre at one end (silver) to ten million billion ohm metres at the other (fused quartz). No linear axis could show all eleven on one page.

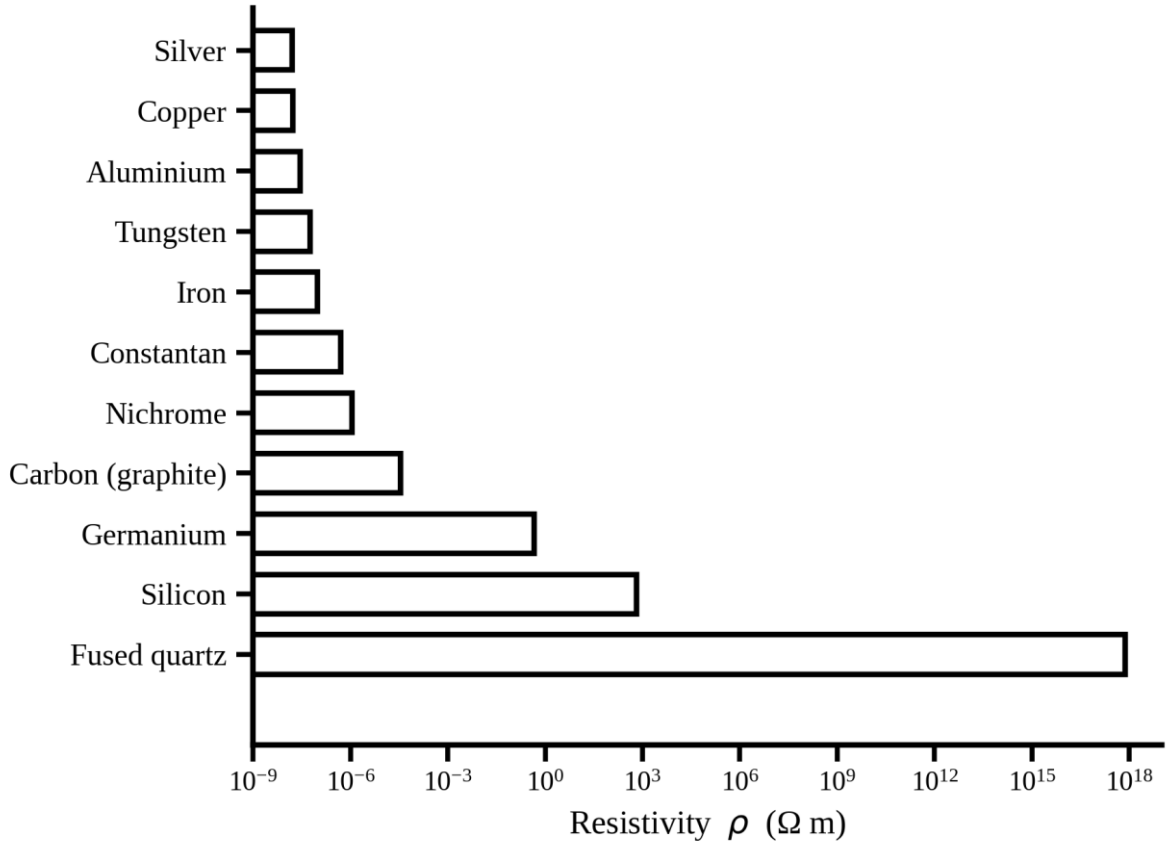


Figure: ***Resistivity of representative materials at 20°C, plotted on a logarithmic axis.*** *The eleven entries span twenty-five orders of magnitude. Silver, copper, aluminium, tungsten, and iron are the conductors. Constantan and nichrome are the designed-inconvenient alloys, used where resistance is wanted rather than avoided. Carbon graphite, germanium, and silicon are semiconductors, sitting between the conductors and the insulators. Fused quartz is an insulator.*

Four classes are visible at a glance.

First, the metallic conductors at the top of the chart (silver, copper, aluminium, tungsten, iron). Their resistivities lie within one order of magnitude of one another and all sit near $10 \times 10^{-9}\Omega$m (ten nanoohm metres). Silver is the best, but copper is cheaper and only seven percent worse, which is why almost every wire in your house is copper. Aluminium is sixty percent worse than copper, but it is also three times lighter, which is why long-distance transmission lines are aluminium and not copper. Tungsten is the metal of choice for filament lamps, not because it conducts best but because it stays solid at the high temperatures the filament must reach to glow visibly.

Second, the resistance-wire alloys (constantan, an alloy of copper and nickel; nichrome, an alloy of nickel and chromium). Their resistivities are about thirty to fifty times that of copper. This is not a flaw of the alloys; it is the whole point of using them. A laboratory rheostat or a kitchen kettle's heating element needs a wire whose resistance is several ohms or tens of ohms in a short length, rather than the thousandth of an ohm copper would give. Constantan and nichrome are designed to resist, not to conduct. They are also chosen because, as we shall see in a moment, their resistance changes very little when the temperature changes.

Third, the semiconductors (carbon graphite, germanium, silicon). Their resistivities lie between those of metals and those of insulators. Pure silicon at room temperature has resistivity about 640Ωm, fifty billion times that of copper. Doped silicon (the

silicon out of which every chip in every phone is made) is many orders of magnitude better, but still nowhere near a metal. The middle of the chart belongs to the semiconductors, and the next book in this series devote a whole chapter to them.

Fourth, the insulators. Fused quartz is the example shown, with resistivity around $10^{17}\,\Omega\text{m}$. Glass, porcelain, dry wood, and most plastics are similar. Their resistivities are so much larger than those of the conductors that, for practical purposes, no current flows in them under any voltage a student is likely to apply to them. The whole of Chapter 11 quietly depended on this fact: the dielectric between the plates of a capacitor was assumed not to conduct at all, even when held at thousands of volts, and the chart shows why.

How resistance changes with temperature

Take any of the metal wires from the chart and warm it. Heat a copper coil from cool tap-water temperature to the temperature of boiling water and its resistance rises by about a quarter. Heat the tungsten filament of an incandescent lamp from the room temperature it sits at when the lamp is off to the 2500K it reaches when the lamp is on, and the resistance rises by a factor of about ten. Heating a metal does not just damage it; it also, and reliably, makes it conduct worse. The microscopic picture tells us at once why.

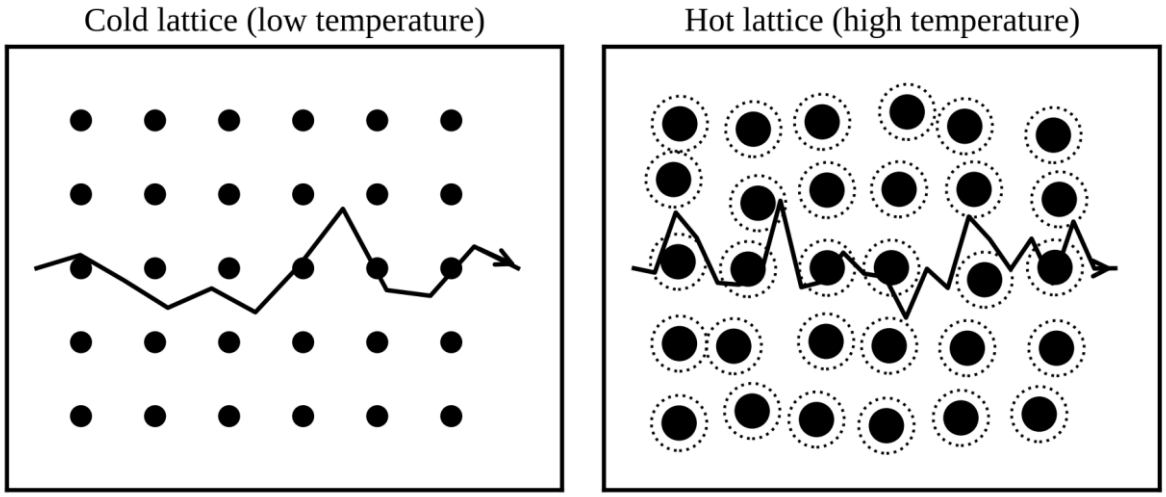


Figure: ***An electron threading through a metal lattice at two temperatures.*** *On the left, a cold lattice: the ions sit almost still and the electron makes its way past them with few collisions. On the right, a hot lattice: the ions vibrate vigorously about their lattice sites, sweeping out the dotted halos shown. The electron now has more frequent collisions, its relaxation time is shorter, and the wire's resistance is larger.*

From the last section, the resistivity is $\rho = \frac{m}{ne^2\tau}$. Of the four factors on the right, three of them (the electron's mass m, the electronic charge e, and the free-electron number density n) are essentially fixed for any given metal; they do not change when the metal is warmed. The fourth factor, the relaxation time τ, does change with temperature, and that is where the whole effect lives. The ions of the lattice vibrate more violently as the metal is heated, and an electron drifting through the lattice has a larger target to collide with at every site it passes. Collisions become more frequent, the relaxation time shortens, the resistivity rises, and (through $R = \frac{\rho l}{A}$) the resistance also rises.

If you do this experiment carefully over a moderate range, say from $0°\text{C}$ to a few hundred degrees, the rise of R with temperature is very close to linear. Write the temperature as θ (in degrees Celsius), call the resistance at the reference temperature θ_0 by the symbol R_0, and call *the fractional change in resistance per unit change in temperature* **coefficient of resistance**, written with the Greek letter $\boldsymbol{\alpha}$. Then, to a good first approximation:

$$\boldsymbol{\Delta R = \alpha R_0 \Delta\theta}$$

Here ΔR is the change in resistance from R_0, and $\Delta\theta$ is the change in temperature from θ_0. The temperature coefficient α has units of per degree Celsius (equivalently per kelvin, since the **size** of a degree Celsius equals the size of a kelvin). For most pure metals α is small and positive, of order a few thousandths per degree Celsius.

Writing $R_\theta - R_0$ in place of ΔR, and $\theta - \theta_0$ in place of $\Delta\theta$, gives:

$$R_\theta - R_0 = \alpha R_0(\theta - \theta_0)$$
$$R_\theta = \alpha R_0(\theta - \theta_0) + R_\theta$$

Hence:

$$R_\theta = R_0(1 + \alpha(\theta - \theta_0))$$

With the reference temperature θ_0 of $0°\text{C}$, the formula simplifies further to:

$$\boldsymbol{R_\theta = R_0(1 + \alpha\theta)}$$

Three things about it must be remembered together, because the formula loses its meaning if **any** of the three is forgotten.

First thing: The reference temperature is $0°C$, not room temperature and not anything else. The value of α you read from a textbook table is the one that makes the formula above work with R_0 taken at exactly $0°C$. Whenever a problem gives you 'the temperature coefficient of copper is $4.3 \times 10^{-3}\ °C^{-1}$', that number is locked to the $0°C$ reference.

Second thing: The temperature θ in the formula is measured in degrees Celsius, not kelvin. The numerical value of α is the same in either scale (because the size of one degree Celsius equals the size of one kelvin), but the value of θ you substitute is the Celsius reading.

Third change: α has units of per degree Celsius (or equivalently per kelvin). For most pure metals α is small and positive, of order a few thousandths per degree Celsius, and a table of representative values is given in the next subsection.

There is one more form worth knowing.

In practice nobody can chill a wire to exactly $0°C$ just to read off R_0. Most laboratory experiments work at convenient temperatures well above zero. So we almost never use the strict formula $R = R_0(1 + \alpha\theta)$ directly. We use, instead, a working form derived from it. Apply the strict formula at two different temperatures θ_1 and θ_2 of the same wire and divide one equation by the other:

$$\frac{R_1}{R_2} = \frac{R_0(1 + \alpha\theta_1)}{R_0(1 + \alpha\theta_2)} = \frac{1 + \alpha\theta_1}{1 + \alpha\theta_2}$$

The R_0 at $0°C$ cancelled out completely. What remains is a clean relation between the two measured resistances, the two known temperatures, and the single material constant α. This is the ratio form of the temperature-coefficient relation, and it is the form that does almost all of the practical work in this section:

$$\boldsymbol{\frac{R_1}{R_2} = \frac{1 + \alpha\theta_1}{1 + \alpha\theta_2}}$$

Two facts about the ratio form are worth stating explicitly. First, it is mathematically exact (no approximation has been made beyond the original linear assumption); it is just the strict formula divided by itself at two temperatures. Second, the reference temperature $0°C$ is still present, but invisibly: the α inside the ratio form is the same α defined at $0°C$, and the table values can be used directly. You never have to measure or know R_0.

A word of warning about the general form:

$$R(\theta_2) = R(\theta_1)[1 + \alpha(\theta_2 - \theta_1)]$$

This form looks like the strict formula with the reference moved from $0°C$ to a more convenient temperature θ_1. Be careful: it is not exact. It is an approximation that treats $R(\theta_1)$ as if it were R_0, which it is not (it is $R_0(1 + \alpha\theta_1)$). The approximation is good only when $\alpha\theta_1$ is much smaller than one. For copper at room temperature $20°C$, $\alpha\theta_1 \approx 0.086$, which is not at all negligible, and the resulting error in the predicted resistance is several percent. Throughout this book we therefore use the ratio form (which is exact) and avoid the approximate general-reference form.

A graph of resistance versus temperature

From the exact formula:

$$R_\theta = R_0(1 + \alpha\theta),$$
$$\boldsymbol{R_\theta = R_0\alpha\theta + R_0}$$

So the graph of resistance R_θ versus temperature θ is a straight line with positive slope $R_0\alpha$ and y-intercept R_0.

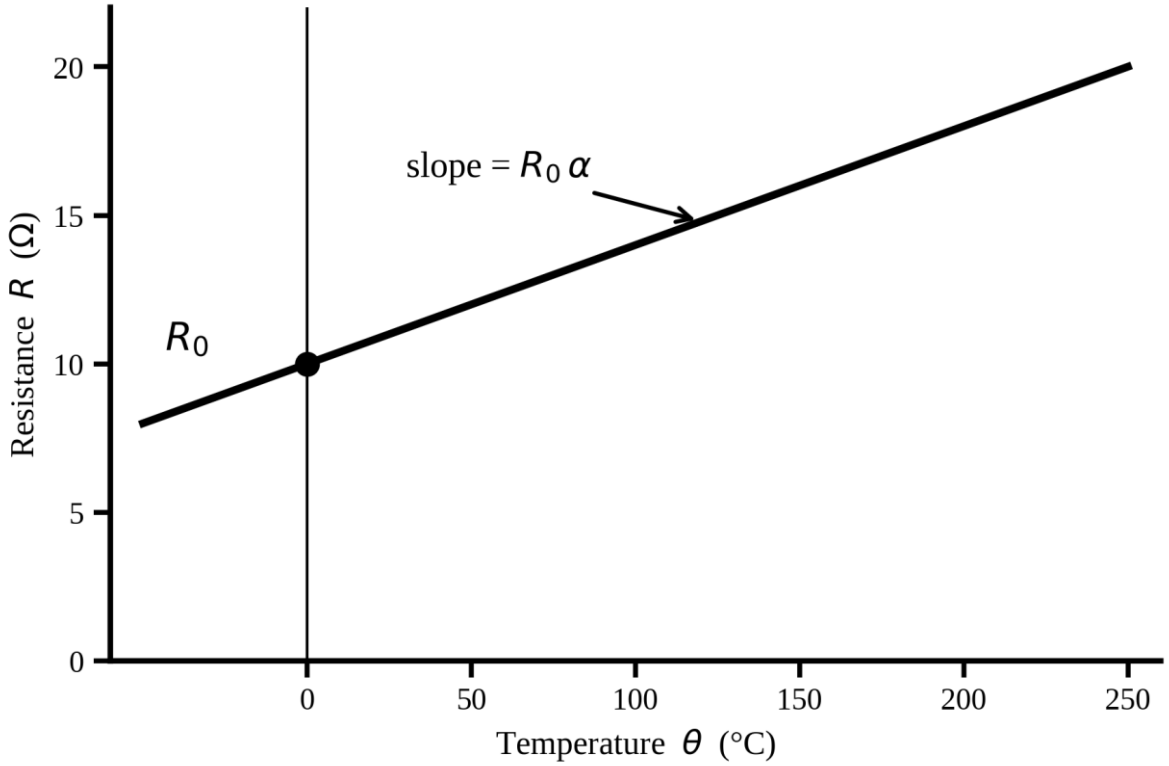


Figure: ***Resistance of a metal as a function of temperature.*** *The plot is a straight line over the range shown, with intercept R_0 at $\theta = 0°C$ and slope $R_0\alpha$. The graph would deviate from a straight line at very low temperatures (near absolute zero) and at very high temperatures (near the melting point); in the working range of any ordinary apparatus it is straight to within a fraction of a percent.*

Some representative values of α:

Material	**Temperature coefficient α at 0°C (per °C)**
Silver	4.0×10^{-3}
Copper	4.3×10^{-3}
Aluminium	4.0×10^{-3}
Tungsten	4.5×10^{-3}
Iron	5.0×10^{-3}
Constantan	4×10^{-4}
Nichrome	4×10^{-4}
Manganin	2×10^{-5}
Pure carbon (graphite)	-5×10^{-4}
Silicon (pure)	-75×10^{-3}

The table holds three lessons in plain view.

The first lesson: Pure metals (copper, aluminium, tungsten, iron, silver) have α values that cluster near 4×10^{-3} per degree Celsius. This is no coincidence; it tracks the way lattice scattering grows with temperature in a metal whose free-electron density is fixed and whose lattice ions vibrate more violently as the metal is warmed. Different metals have different ion masses and different lattice spacings, so the values are not identical, but they are all of the same order.

The second lesson: The alloys constantan and nichrome, and the better alloy manganin, have α values that are smaller than copper's by roughly a factor of ten to a thousand. Their resistance changes only a little when their temperature changes, which

is exactly the property a standard laboratory resistance coil or a heating element needs. The alloys are designed for that purpose, and the design is reflected in their small α.

The third lesson: Not every α is positive. Pure carbon graphite and silicon both have negative temperature coefficients. Their resistance falls when they are heated. The microscopic reason in these substances is that the number density of carriers, n, is not constant with temperature as it is in metals. Heating a semiconductor frees more electrons into the conduction band, so n rises with temperature. The relaxation time τ still shortens with temperature, but the rise in n wins, and the net effect is that the resistance falls.

What α tells you about the metal

A large positive α (such as iron's, 5×10^{-3} per degree Celsius, or copper's, 4.3×10^{-3} per degree Celsius) means the lattice scatters drifting electrons strongly, and the strength of the scattering grows quickly with temperature. Copper between $0°C$ and $100°C$ raises its resistance by 43 percent; the same coil between $0°C$ and $50°C$ raises it by 22 percent. If you wanted to design an electric thermometer that worked by reading the resistance of a coil, copper would be a fine choice: its α is large enough that a modest temperature change gives a clearly measurable resistance change.

A small positive α (constantan's, nichrome's, or manganin's, of order 10^{-4} per degree Celsius) means the lattice scattering is nearly independent of temperature in the working range. Coils made of these alloys keep their resistance almost unchanged as their temperature changes. That is why every good-quality laboratory resistor in your physics laboratory is made of nichrome or constantan and not of copper, why the heating element of a kettle is nichrome, and why the same answers come out of the same resistance experiment whether the laboratory is cool at dawn or warm by mid-afternoon.

A negative α (carbon graphite, silicon, germanium) means the substance behaves opposite to a metal. Heat the substance, and its resistance falls. The microscopic mechanism is different: in these substances the number of carriers per cubic metre is itself a function of temperature, rising as the substance is heated, and that rise dominates everything else. We will meet the consequences again in the chapters that follow, where a gas at low pressure passes from being a perfect insulator to being a strong conductor once enough energy is poured into it to ionise its molecules.

How resistance changes with the geometry of the wire

We separated the resistance into a material part (ρ) and a geometric part $\left(\frac{l}{A}\right)$. The material part has now been fully discussed. The geometric part needs three short comments before we are done with this section.

Case one is the easy case. Hold the cross-sectional area fixed and change only the length. Then $R = \frac{\rho l}{A}$ and R is proportional to l. Doubling the length doubles the resistance. Halving the length halves it. Nothing surprising.

Case two is more delicate. Take a wire of length l and cross-section A, and stretch it (without breaking it) to a longer length while letting it thin out. The substance is incompressible, so the volume $V = Al$ is preserved as the wire is stretched. If the length becomes k times its original value, the cross-section must become $\frac{A}{k}$, so that the volume is still Al. The new resistance is:

$$\mathbf{R_{new}} = \frac{\rho(kl)}{\frac{A}{k}} = k^2 \frac{\rho l}{A} = \mathbf{k^2 R}$$

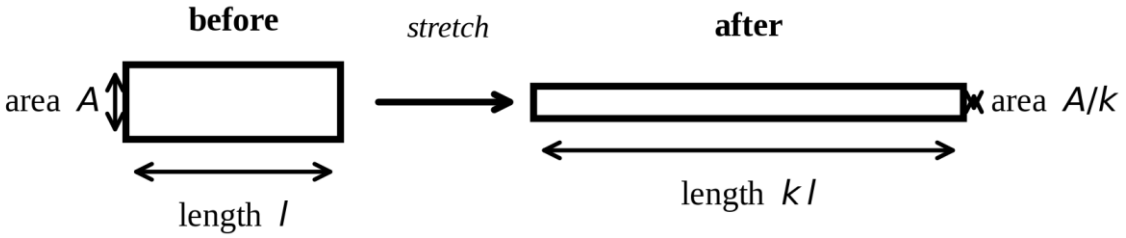


Figure: ***A wire stretched at constant volume.*** *The length grows from l to k l, the cross-sectional area shrinks from A to A divided by k, and the resistance grows by a factor of k squared, not just k.*

Stretching the wire to twice its length, at constant volume, multiplies the resistance by four, not by two. The stretching is doing two things to the resistance at once: *it is lengthening the path the electrons must travel,* ***and*** *it is narrowing the path they must travel through.* Both effects raise R by the same factor k, so the combined effect is $\mathbf{k^2}$. Forgetting the second of the two effects is the most common error students make with the stretched-wire question, and the worked Example 15 catches it directly.

Case three is fully general. The length and the cross-section are changed independently, perhaps because the wire is being redrawn through a die that gives it a different diameter as well as a different length. Then $R = \frac{\rho l}{A}$ reads off directly: if l becomes k times itself and A becomes j times itself, the new resistance is:

$$R_{new} = \frac{k}{j} R$$

with no constraint between k and j, because the volume is no longer being held fixed. The constant-volume case ($kj = 1$) recovers the stretched-wire result of case two, and the constant-area case ($j = 1$) recovers case one.

Non-ohmic conductors and their I-V characteristics

The last section made the point in passing, and the present section has now given us the tools to look properly at the non-ohmic conductors that section postponed. The figure below shows the current-voltage characteristics of four conductors on common axes.

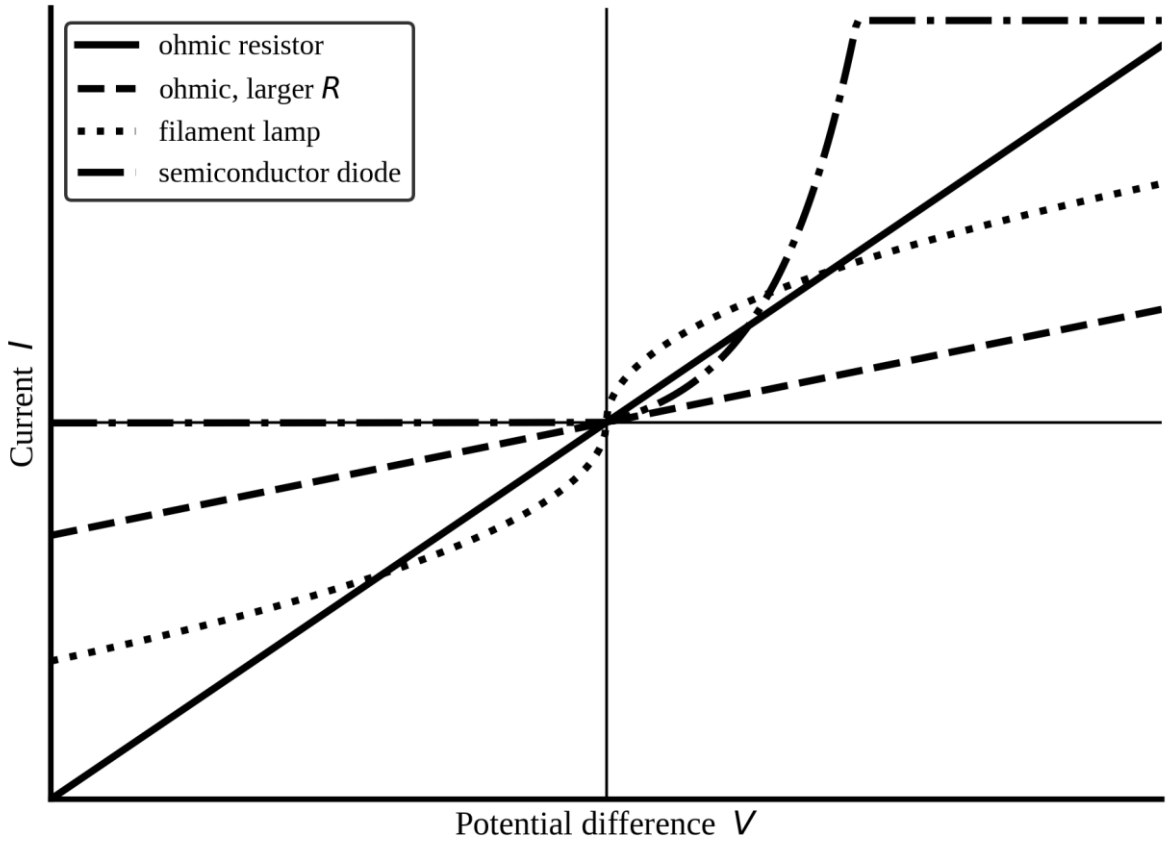


Figure: ***Current-voltage characteristics of four conductors.*** *Solid line: an ordinary ohmic resistor. Dashed line: a second ohmic conductor with larger resistance (smaller slope). Dotted line: a filament lamp, whose resistance rises as the filament heats with rising current. Dash-dot line: a semiconductor diode, which conducts easily in one polarity and almost not at all in the other.*

The first two conductors are ohmic, and their curves are straight lines through the origin with different slopes. The reciprocal of the slope is the resistance, so the steeper line belongs to the smaller resistance. Reversing the polarity of V reverses the polarity of I; the line is straight on both sides of the origin. Nothing here is new.

The dotted curve is the filament lamp. At very small V, the filament is cool, its resistance is small (about ten times smaller than its hot value), and the slope of the curve is correspondingly steep. As V is raised, the current rises, the I^2R heating raises the filament's temperature, the filament's resistance rises (because tungsten has a large positive α, 4.5×10^{-3} per degree Celsius), and the slope of the curve falls. By the time the filament is glowing at $2500K$ its resistance is roughly ten times the room-temperature value, and the curve has bent over noticeably. The filament is not disobeying Ohm's law at any moment; at any moment its resistance is well defined, and the current obeys $I = \frac{V}{R}$ at that resistance. But the resistance is changing as the filament heats up. The dotted curve is the trace of a moving Ohm's law, not the failure of one.

The dash-dot curve is the semiconductor diode. The diode's behaviour is asymmetric in polarity. With one polarity (forward bias) the diode passes current that rises exponentially with voltage; with the other (reverse bias) the diode passes a tiny saturation current that hardly grows at all until the diode is destroyed by breakdown at some large reverse voltage. The

asymmetry has no analogue in any ohmic conductor; the device treats positive and negative voltages differently because its internal structure (the junction between two regions of doped silicon) breaks the symmetry. The physics behind the diode is the subject of a later chapter; for now the curve simply confirms that non-ohmic conduction is possible, common, and useful.

Seven examples now follow.

BINDER Example 10

A wire taken from a science-kit drawer at Miono Secondary School has length 10m and a uniform cross-sectional area of $1 \times 10^{-7}\text{m}^2$. A digital multimeter measures its resistance between the two ends as 50Ω at the laboratory temperature of $20°\text{C}$. Compute the resistivity of the wire material and identify the material by comparison with the bar chart of resistivities you read earlier in this section.

Solution

Take the geometric form of the resistance, $\text{R} = \frac{\rho \text{l}}{\text{A}}$, and rearrange it for the resistivity.

$$\rho = \frac{\text{RA}}{\text{l}}$$

Substituting the three measured values:

$$\rho = \frac{(50\Omega)(1 \times 10^{-7}\text{m}^2)}{(10\text{m})} = 5 \times 10^{-7}\Omega\text{m}$$

Compare this value with the bar chart of resistivities given earlier in this section. The closest entry is constantan, whose resistivity is $4.9 \times 10^{-7}\Omega\text{m}$ at $20°\text{C}$, identical to within the precision of the data. The wire is therefore constantan.

Making Sense of the Answer: *The number is large, by metal's standards. It is thirty times larger than copper's resistivity and ten times larger than aluminium's. That is exactly why constantan exists. A coil of constantan wire only ten metres long with a thin one-tenth-of-a-square-millimetre cross-section gives a useful resistance of* 50Ω. *The same coil made of copper, with the same length and the same cross-section, would give* 1.7Ω, *which is hardly noticeable in any circuit. The designed inconvenience of the alloy is what makes laboratory rheostats and standard resistors possible.*

Think Like a Physicist: *In identification problems of this kind, never compare the resistance of two objects. The resistance depends on the shape; two different samples of the same substance can have wildly different resistances. Always compute the resistivity first. The resistivity is a property of the substance, and that is what you compare against the table.*

BINDER Example 11

A short coil of metal wire has resistance 30Ω at $15°\text{C}$ and 36Ω at $65°\text{C}$. Find the temperature coefficient of resistance of the wire material, and use the value to suggest which of the metals in the table the wire is most likely made of.

Solution

The two resistance-temperature pairs lie on the same straight line $\text{R} = \text{R}_0(1 + \alpha\theta)$, but the value of R_0 is not given. The form to use is the ratio form, which has eliminated R_0 already:

$$\frac{\text{R}_1}{\text{R}_2} = \frac{1 + \alpha\theta_1}{1 + \alpha\theta_2}$$

Substitute $\text{R}_1 = 30\Omega$, $\text{R}_2 = 36\Omega$, $\theta_1 = 15°\text{C}$, $\theta_2 = 65°\text{C}$, and rearrange:

$$\frac{30}{36} = \frac{1 + 15\alpha}{1 + 65\alpha}$$

Solve for α:

$$\alpha = 4.26 \times 10^{-3}\ °\text{C}^{-1}$$

Compare this value with the table of α values given earlier in this section. The closest match is copper, whose temperature coefficient is 4.3×10^{-3} per degree Celsius. The wire is therefore very probably copper.

Making Sense of the Answer: *The answer is a single physical number with a clear physical meaning. An* α *of about* 4.3×10^{-3} *per degree Celsius says that every one-degree rise of the wire's temperature raises its resistance by 0.43 percent. Across the* $50°C$ *span of the experiment, the resistance rose by exactly 0.43 percent times 50, which is 21.5 percent. The actual rise was from* 30Ω *to* 36Ω, *which is 20 percent, and the small discrepancy is the result of taking* $15°C$ *rather than* $0°C$ *as the lower endpoint. The numbers are consistent with each other.*

Think Like a Physicist: *When you have two resistances at two temperatures and no third piece of information, the ratio form of the temperature coefficient formula is the form to use. It has already done the work of eliminating the unknown* R_0. *Trying to use the single-temperature form would force you to invent a value of* R_0, *propagate it through the calculation, and then (if you were lucky) cancel it at the end. The ratio form is doing exactly this cancelling for you, only quicker.*

BINDER Example 12

The Bagamoyo agricultural college is rewinding the field coil of an irrigation-pump motor, using bare copper wire. After the rewinding is complete and the motor is still cold, a workshop technician measures the resistance of the new coil at the morning

temperature of $20°C$ and reads 45Ω. Under load on a hot afternoon the coil reaches a steady temperature of $75°C$. Predict the resistance of the coil at that operating temperature. Take the temperature coefficient of copper to be $4.3 \times 10^{-3}°C^{-1}$.

Solution

Again the ratio form is the natural tool: it relates two resistances at two temperatures using only the table α (which is locked to the $0°C$ reference), and it never requires us to know the wire's resistance at $0°C$:

$$\frac{R_{75}}{R_{20}} = \frac{1+75\alpha}{1+20\alpha}$$

$$R_{75} = \left(\frac{1+75\alpha}{1+20\alpha}\right) R_{20} = \left(\frac{1+75\times 4.3\times 10^{-3}}{1+20\times 4.3\times 10^{-3}}\right) \times 45\Omega = 54.8\Omega$$

Making Sense of the Answer: *The resistance of the operating coil is 54.8Ω, about 22 percent larger than the cold value of 45Ω. That is a substantial change, and it has real consequences for the pump. The current drawn from the supply (and hence the mechanical torque the coil can develop) is $I = \frac{V}{R}$, so the same supply voltage delivers about 18 percent less current to the hot coil than to the cold one (since $1/1.218 \approx 0.82$). Engineers who design winding-based motors must allow for this loss of grip when the motor is hot, by increasing the supply voltage or by accepting that the cold machine is the strong one.*

Think Like a Physicist: *The ratio form does two useful things at once. It eliminates the unknown R_0 (the resistance at $0°C$, which the technician would never measure in practice), and it keeps α honest with the table convention. The approximate general-reference form $R(\theta_2) = R(\theta_1)[1+\alpha(\theta_2-\theta_1)]$ would have given 55.6Ω here, an answer wrong by about 1.5 percent. Small, but real, and a sign that the wrong form was used. Whenever you find yourself with two temperatures and an α from a table, reach for the ratio form first.*

HOT Example 13

The nichrome heating element of a Mbeya bakery oven has resistance 20Ω at the cool temperature of $25°C$. When the oven is pre-heated to $125°C$ the resistance of the element measures 22Ω. Find the temperature coefficient of the nichrome, and then predict the resistance of the element when the oven is operating at its full bread-baking temperature of $625°C$.

Solution

The first step is to find α from the two known resistance-temperature pairs.

$$\frac{R_1}{R_2} = \frac{1+\alpha\theta_1}{1+\alpha\theta_2}$$

Substitute $R_1 = 20\Omega$ at $\theta_1 = 25°C$ and $R_2 = 22\Omega$ at $\theta_2 = 125°C$:

$$\frac{20}{22} = \frac{1+25\alpha}{1+125\alpha}$$

$$\alpha = 1.03\times 10^{-3}\ °C^{-1}$$

The value $1.03 \times 10^{-3}°C^{-1}$ is roughly two to three times the value of pure nichrome listed in the table ($4 \times 10^{-4}\ °C^{-1}$). The discrepancy is characteristic of commercial heater wire, which is rarely pure nichrome but is usually a nichrome-based alloy with small amounts of other metals added to make it easier to draw and to weld. The value to use for this particular element, calibrated from this experiment, is the experimentally determined $1.03 \times 10^{-3}\ °C^{-1}$.

With α now known, apply the ratio form once more to find the resistance at $625°C$. Pair the unknown R_{625} at $625°C$ with the known 20Ω at $25°C$:

$$\frac{R_{625}}{R_{25}} = \frac{1+625\alpha}{1+25\alpha}$$

$$R_{625} = \left(\frac{1+625\alpha}{1+25\alpha}\right) R_{25} = \left(\frac{1+625\times 1.03\times 10^{-3}}{1+25\times 1.03\times 10^{-3}}\right) \times 20\Omega = 32\Omega$$

Making Sense of the Answer: *The element's resistance has risen by 60 percent between the cool oven and the hot oven. The current it draws from the mains, at constant voltage, has fallen by the same fraction (roughly speaking; in fact the mains voltage in Tanzania, $240V$, is held very steady by the grid). The element therefore self-regulates its temperature once it reaches operating heat: any further rise in temperature would raise the resistance, drop the current, and reduce the power dissipated, which would let the element cool back down. The Mbeya baker is benefiting from this self-regulation every time the oven holds a steady temperature without an explicit thermostat.*

Think Like a Physicist: *In any problem with three temperatures and two unknown quantities (here α and the third resistance), and no resistance value given at the reference temperature $0°C$, the ratio form is the form to use. It eliminates R_0 cleanly in the calibration step (where two known pairs determine α), and it can be used again in the prediction step (where the now-known α and one of the original known pairs together determine the third resistance). Same equation, two substitutions, no fictitious reference value invented along the way. This is the lesson Example 11 made explicit, and it applies here word for word.*

HOT Example 14

A piece of copper wire has resistance 8Ω. The wire is melted down and redrawn through a die so that it emerges four times longer than before and three times thicker in diameter. Compute the resistance of the redrawn wire.

Solution

The wire is no longer the same wire; it has been remade. But it is still copper, so the resistivity ρ is unchanged. The length has been multiplied by $k = 4$, and the diameter has been multiplied by 3, which multiplies the cross-sectional area (proportional to the square of the diameter) by $j = 9$. From the geometric form of the resistance:

$$R = \frac{\rho l}{A}, \quad R_{new} = \frac{\rho(kl)}{jA} = \frac{k}{j}R$$

Substituting $k = 4$, $j = 9$, $R = 8\Omega$:

$$R_{new} = \frac{4}{9} \times 8\Omega = 3.56\Omega$$

Making Sense of the Answer: *The lengthening tries to raise the resistance by a factor of 4, but the thickening tries to lower it by a factor of 9, and the thickening wins. The redrawn wire has less than half the resistance of the original, even though the new wire is four times the length. Doubling the diameter of a wire is far more effective at lowering its resistance than is shortening the wire: the area depends on the square of the diameter, and the resistance depends inversely on the area.*

Think Like a Physicist: *Always work out separately what each geometric change does, then combine them. Length and area enter R differently: length in the numerator, area in the denominator. Mixing them up is the most common error in stretched-wire questions, and the remedy is the formula* $R = \frac{\rho l}{A}$ *written out in full, with every factor visible.*

HOT Example 15

A uniform wire of length 30cm has resistance 0.40Ω. The wire is then carefully stretched, without breaking, to a new length of 45cm. Throughout the stretching the wire keeps the same total volume of material. Compute the resistance of the stretched wire.

Solution

Two things happen to the wire when it is stretched at constant volume. The length grows; the cross-section shrinks. The ratio between the new length and the old length is:

$$k = \frac{l_{new}}{l_{old}} = \frac{45\text{cm}}{30\text{cm}} = 1.5$$

Because the volume is preserved, the cross-section shrinks by exactly the same factor, so the new area is $\frac{A}{k}$. From the geometric form of the resistance:

$$R_{new} = \frac{\rho(kl)}{\frac{A}{k}} = k^2\frac{\rho l}{A} = k^2R$$

Substituting $k = 1.5$ and $R = 0.40\Omega$:

$$R_{new} = (1.5)^2 \times (0.40\Omega) = 0.90\Omega$$

Making Sense of the Answer: *The resistance has more than doubled, even though the length has only grown by half. This is the classic stretched-wire result: at constant volume, a factor of k stretching of the length multiplies the resistance by* k^2*, not by k. Stretching the wire is doing two harmful things to the current path at the same time, lengthening it and narrowing it, and the two effects compound multiplicatively, not additively.*

Think Like a Physicist: *This is one of those problems that almost answers itself once you notice the constant-volume constraint. The whole answer lives in the substitution A becomes* $\frac{A}{k}$*. If the wire instead were stretched at constant cross-section (impossible in practice, but worth thinking about), the resistance would only have risen by a factor of 1.5. The compounding factor of* k^2 *is paid for by the conservation of metal.*

REAL Example 16

TANESCO is planning a 280km transmission line from Mwanza to Kahama. The engineering committee is choosing between two metals for the line, copper and aluminium. The line must have the same total resistance per kilometre regardless of which metal is chosen. For each metal, find the cross-section of the conductor required, and compare the total mass of the two candidate lines.

TANESCO is planning a 280km transmission line from Mwanza to Kahama. The engineering committee is choosing between copper and aluminium conductors. The line must be designed so that each conductor has a resistance of 0.10Ω per kilometre at $20°\text{C}$.

(a) Determine the cross-sectional area required for each metal.
(b) Hence calculate the total mass of the complete 280km line for each metal.

(c) Compare the masses of the two lines and comment on which material is more practical for long-distance transmission.

Take the resistivities at $20°\text{C}$ to be $1.7 \times 10^{-8}\Omega\text{m}$ for copper and $2.8 \times 10^{-8}\Omega\text{m}$ for aluminium.

Take the densities to be 8900kg/m^3 for copper and 2700kg/m^3 for aluminium.

Solution

First, find the cross-section required for each metal. The resistance of one kilometre of conductor is $\frac{\rho l}{A}$. Setting this equal to the specified $0.10\Omega/\text{km}$ and rearranging:

$$A = \frac{\rho l}{R}$$

For copper, with $\rho = 1.7 \times 10^{-8}\Omega\text{m}$, $l = 1000\text{m}$, $R = 0.10\Omega$:

$$A_{Cu} = \frac{(1.7 \times 10^{-8}\Omega\text{m})(1000\text{m})}{(0.10\Omega)}$$

$$A_{Cu} = 1.7 \times 10^{-4}\text{m}^2$$

For aluminium, with $\rho = 2.8 \times 10^{-8}\Omega\text{m}$, the same length, the same required resistance:

$$A_{Al} = \frac{(2.8 \times 10^{-8}\Omega\text{m})(1000\text{m})}{(0.10\Omega)}$$

$$A_{Al} = 2.8 \times 10^{-4}\text{m}^2$$

Now, for each metal, the mass of the full 280km line is given by:

$$m = \rho V = \rho A L$$

where ρ is the density of the metal and L is the total length of the line, $2.80 \times 10^5\text{m}$. For copper:

$$m_{Cu} = (8900\text{kg/m}^3)(1.7 \times 10^{-4}\text{m}^2)(2.80 \times 10^5\text{m}) = 4.24 \times 10^5\text{kg}$$

For aluminium:

$$m_{Al} = (2700\text{kg/m}^3)(2.8 \times 10^{-4}\text{m}^2)(2.80 \times 10^5\text{m}) = 2.12 \times 10^5\text{kg}$$

The mass ratio is therefore:

$$\frac{m_{Al}}{m_{Cu}} = \frac{2.12 \times 10^5}{4.24 \times 10^5} = 0.50$$

The aluminium line is half the mass of the copper line, even though it is two-thirds thicker. The mass saving wins on every count: every pylon must support a lighter cable, every footing must hold up a lighter pylon, and every truckload of wire delivered to the construction site costs half as much freight. Aluminium is therefore the metal TANESCO chooses for Mwanza-Kahama, as it does for every long-distance transmission line in the country.

Making Sense of the Answer: *The mass ratio of one-half is no accident. The ratio of the resistivities is roughly $\frac{2.8}{1.7}$, and the ratio of the densities is roughly $\frac{2.7}{8.9}$. The product of the two ratios is $\frac{2.8\times2.7}{1.7\times8.9}$, which works out to 0.50. The aluminium line is a little less efficient electrically (per cubic metre) than the copper line, but it compensates by being more than three times lighter (per cubic metre), and the compensation more than wins. Every transmission line in the world reflects the same compromise.*

Think Like a Physicist: *Engineering design is always a multi-criterion optimisation, and the right answer is the one that matters most for the application. For overhead transmission, mass and cost beat conductivity, and aluminium wins. For the wiring of a household, where the runs are short and the cross-section is constrained by the wall conduit, conductivity beats mass, and copper wins. The physics is the same in both cases (the formula $R = \frac{\rho l}{A}$, and the densities of the two metals); the choice changes with the context. A good physicist always asks not just for the answer but for the constraints that made the answer matter.*

Seven examples behind us. We have identified a science-kit wire by its resistivity, ambushed a coil into confessing its temperature coefficient, predicted what a Bagamoyo motor will do when the afternoon turns hot, watched a Mbeya oven calibrate itself by the very rise in resistance that the baker once cursed, made a copper wire shrink its own resistance by getting fatter faster than it got longer, made another wire double its resistance simply by agreeing to stretch, and walked a TANESCO committee through the single piece of arithmetic that decides whether 280km of transmission line are made of copper or of aluminium. None of those answers was a slogan; every one of them came out of a single formula, $R = \frac{\rho l}{A}$, combined with the way ρ itself responds to the temperature of the wire. The chapter so far has given you everything you need to read the resistance of any real conductor as a story about its material and its geometry, and to predict what that story will say next when the conductor is warmed, cooled, stretched, or remade.

With the resistance of the wire itself now thoroughly understood, one question remains, and it is the question we have been carefully sidestepping since the first paragraph of this chapter. Where does the current come from in the first place? Every example so far has begun with the words ‘a battery of voltage V is connected to…’, but no description of the battery itself has been offered. *What is inside a battery? What does it actually do to the charges in the wire? Why does the voltage across a battery’s terminals fall a little when current is drawn from it, and rise back when current is stopped? Is there a hidden*

resistance inside every battery that no manufacturer advertises? The next section answers all four questions, and the key idea is one that every dry cell, every car battery, every solar inverter, and every transmission-line generator have in common: **the electromotive force**.

ELECTROMOTIVE FORCE, INTERNAL RESISTANCE, AND TERMINAL VOLTAGE

The closing paragraph of the last section asked what is inside a battery, what it actually does to the charges flowing through a wire, and why the voltage across a battery's terminals falls a little when current is drawn from it. All three questions can be answered together with a single piece of vocabulary and a single short formula. The vocabulary comes first.

The electromotive force (EMF) defined

The electromotive force of a source *is the energy that the source supplies to every unit of charge that passes through it.* This is the plain-English definition, and it is worth saying twice in the two ways students most often need to hear it.

- Per coulomb of charge that the source pushes around the circuit, the EMF tells you how many joules of energy the source itself contributes.
- Equivalently, the EMF is the amount of work the source performs on each coulomb of charge in lifting that charge from its lower-potential terminal to its higher-potential terminal, ready to be sent into the external circuit.

A short analogy makes the definition vivid.

Imagine a battery as a tiny pump that lives inside its casing. The pump receives charges arriving at one of its terminals (the negative one), lifts each of them up through some internal step of energy, and releases them at the other terminal (the positive one) with more energy than they started with. The energy gained by each unit of charge during this lift is the **electromotive force** of the battery. The battery acts on charges in roughly the way a hand-pump acts on water: doing work on each unit that passes through, regardless of how much water (charge) passes through in total.

To turn the definition into a formula, suppose the source pumps a total charge Q through itself while doing total work W on those charges. By definition, the **EMF E** is the work done per unit charge:

$$\mathbf{E = \frac{W}{Q}}$$

The units of E follow at once: *joules per coulomb*, which is exactly the **volt.** (One volt is one joule per coulomb.) So when a manufacturer labels a battery 1.5V, the label is telling you that the battery supplies one and a half joules of energy to every coulomb of charge that passes through it.

The name electromotive force is a historical inconvenience and must not be taken literally. Despite the word 'force', the EMF is not a mechanical force. It is not measured in newtons. It has the units of energy per unit charge (joules per coulomb, i.e. volts), just like any other voltage. The word 'electromotive' captures the idea that the EMF is what sets charges into motion in the first place, but the word 'force' is a translation error from older European physics literature that has stuck for two centuries. Whenever you see EMF, mentally substitute 'driving voltage of the source'; that is closer to what it actually means.

Two things are worth saying about E before we go further.

First thing: *E is a property of the source alone, and not of the circuit the source is plugged into.* A AA battery and a car battery, both built from cells of the same chemistry, have nearly the same EMF per cell (about 1.5V for a dry cell, about 2V for each cell in a lead-acid car battery). What differs between them is the current the larger battery can supply, not the EMF. A car battery can deliver hundreds of amperes briefly; a torch cell would melt under the same load. But the energy delivered per coulomb is the same.

Second thing: *E is in general different from the voltage you would measure across the battery's terminals when the battery is in use*. That measured voltage has its own name and its own definition, given in a moment. The two voltages are equal only when no current is flowing through the battery. As soon as current flows, the measured voltage drops below the EMF by a small but definite amount. The reason has a name too, and that name is the second piece of vocabulary for this section.

Internal resistance demystified

The internal resistance of a source, written **r**, *is the total ohmic resistance offered by everything inside the source's casing to the current that the source drives.* The source is not an idealised pump operating in vacuum. Inside its casing live three real, physical things, and each of them resists the flow of current in the ordinary way.

First, the chemistry that does the lifting. Inside any wet cell, charge is carried partly by ions moving through a liquid or paste electrolyte, and the ions experience drag from the surrounding molecules of the electrolyte. That drag is, electrically speaking, an ohmic resistance. **Second,** the connectors that join the chemistry to the external terminals. These are made of metal, but the metal is neither infinitely thick nor perfectly conducting. **Third,** short lengths of wire inside the casing. Each of the three contributes its own small resistance. The sum of all three contributions is the internal resistance of the cell.

The internal resistance behaves like any other ohmic resistor. When the cell drives a current I through its own internal circuitry, it dissipates power I^2r as heat inside itself, exactly the formula for power dissipated in any resistor of resistance r carrying

current I. That heat never reaches the external circuit; it warms the inside of the cell instead. If you have ever held a torch on for ten minutes and then felt the cells warm to the touch, you have felt this dissipation directly. The I^2r warming is the price the cell pays for being a real object rather than an idealised one.

By convention, *r is drawn as a small **zigzag** inside the rectangle that represents the cell*, rather than as a separate resistor outside. The physics is the same either way; the symbol is a shorthand for 'this resistance lives inside the casing'. When you solve a circuit, treat r as you would any other resistor, with the single understanding that it sits in **series** with the EMF source and inside the same casing as the source.

Terminal voltage and the V = E − Ir relation

The **terminal voltage** of a source, written **V**, *is the voltage that the external circuit sees across the source's two terminals while the source is in use.* Operationally, the terminal voltage is what a voltmeter would read if connected directly across the source's two terminals at the moment of measurement. The terminal voltage and the EMF are two different quantities, and one of the main tasks of this subsection is to derive the precise relationship between them.

Suppose the cell drives a steady current I through an external resistor R, as in the figure below.

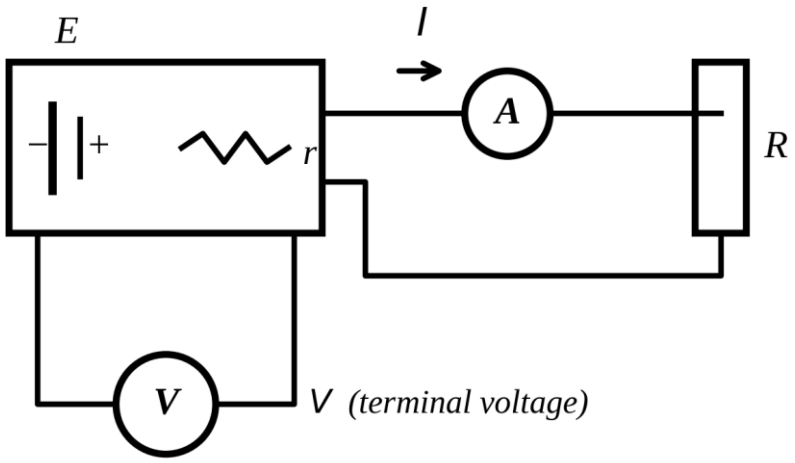


Figure: ***A battery of EMF E and internal resistance r driving a current I through an external resistor R****. The voltmeter, drawn below the battery, reads the terminal voltage V, which is what the external circuit sees. The ammeter in the top wire reads I.*

Trace one coulomb of charge round the circuit, and apply conservation of energy to it. Inside the cell, the chemistry does work E on the coulomb (raising it from the negative terminal up to the positive terminal). Some of that work, of magnitude Ir, is then immediately dissipated as the coulomb pushes through the internal resistance r on its way out. The remainder, $E - Ir$, is what the coulomb actually carries with it into the external circuit. By the operational definition above, that remainder is the terminal voltage of the cell:

$$\mathbf{V = E - Ir}$$

This is the central equation of the section. It says that *the voltage available to the external circuit equals the EMF minus the drop across the internal resistance.* The external circuit never sees the EMF directly; what it sees is the EMF reduced by the internal-resistance loss.

A circuit is said to be in **open circuit** *when no current flows in it*, because the loop is broken somewhere (a switch is open, a wire is cut, or no external resistor is connected). In open circuit, $I = 0$, and the Ir loss in the central equation vanishes. The terminal voltage then equals the EMF exactly:

$$\mathbf{V = E} \quad (\textbf{open circuit})$$

Connect a voltmeter across an open battery, in other words, and the voltmeter reads the EMF directly. This is the cleanest way to measure E in the laboratory; the next figure shows how it is done.

More generally, when r is small and I is moderate, the Ir loss is small and V is only slightly less than E. A good-quality dry cell has r of the order of one-tenth of an ohm; drawing one ampere from such a cell drops the terminal voltage by only 0.1V below the 1.5V EMF, which is hardly noticeable. A worn or cheap cell has a much larger r, which is why an old torch goes dim long before its EMF has decayed.

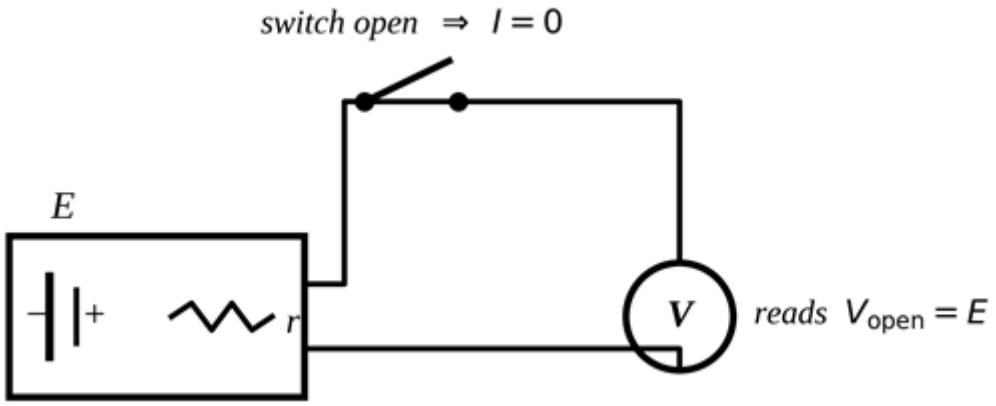


Figure: ***Measuring the EMF of a cell directly.*** *The switch in the top wire is open, so no current flows. The voltmeter then reads V-open, which equals E exactly. In practice the voltmeter must be of very high impedance (defined in the next subsection) so that the tiny current it draws is negligible.*

Plotting V against I for a real source gives a straight line, with two pieces of information on its face. The y-intercept is E (the value of V when $I = 0$), and the slope is $-r$ (the terminal voltage falls by r for every ampere of current drawn). A single experiment that varies the load and records the resulting current and terminal voltage therefore reveals both E and r at once.

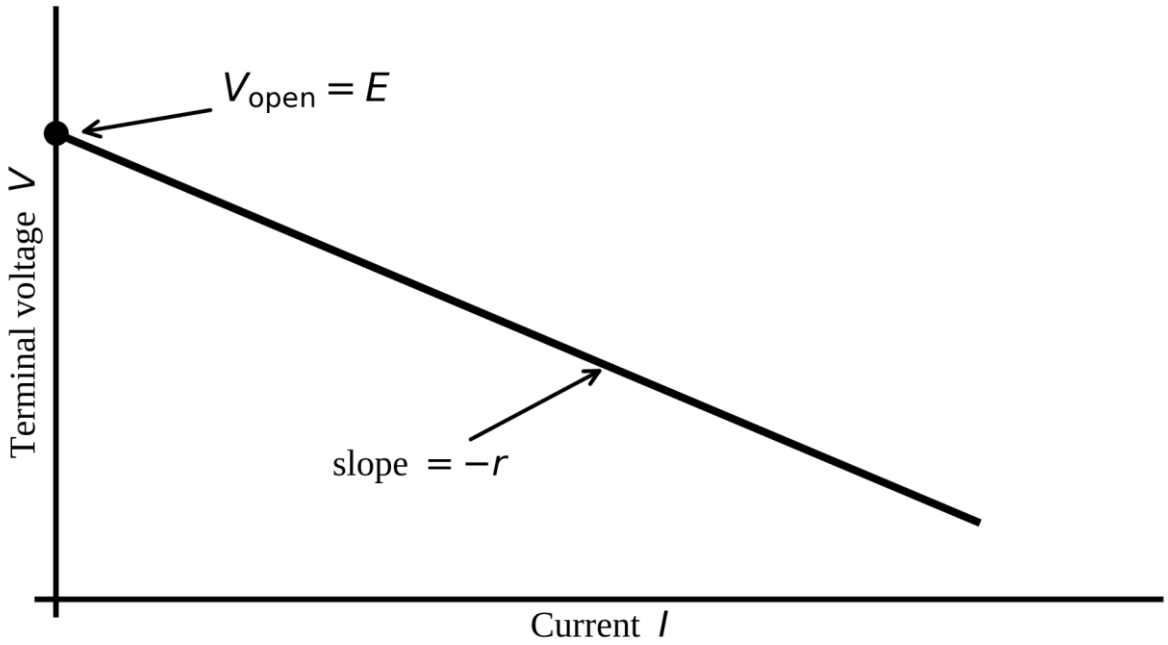


Figure: Terminal voltage as a function of the current drawn from a real source. The line is straight; the y-intercept is the EMF E; the negative slope has magnitude equal to the internal resistance r. The line crosses the I-axis at I = E/r, which is the short-circuit current.

The x-intercept of the line in the figure above, where $V = 0$, introduces another piece of vocabulary, short-circuit current. **The short-circuit current** of a source *is the current that would flow through it if its two terminals were directly joined by a wire of negligible resistance, so that the external resistance R is effectively zero.* Substituting $V = IR = 0$ into the central equation $V = E - Ir$ gives:

$$\mathbf{I_{short} = \frac{E}{r}}$$

The short-circuit current is determined by the EMF and the internal resistance alone, with no external resistance to limit it. It is the **largest current** the source can deliver, and for most sources it is large enough to be dangerous. A car battery (EMF 12V, r of order 0.01Ω) has a short-circuit current of about 1200A, enough to melt a spanner accidentally dropped across its terminals. Never short-circuit a real source deliberately.

Can the terminal voltage ever exceed the EMF? Yes, and the answer introduces one final piece of vocabulary, charging. A source is said to be **charging** *when another, larger source is connected to it in such a way that current is forced backwards through it, into its positive terminal and out of its negative terminal (the opposite of the direction it would drive on its own).*

Charging is what happens to a car battery while the engine is running and the alternator is supplying current, or to a rechargeable phone cell while it is plugged into a wall charger.

During charging the current direction through the internal resistance is reversed, so the sign of the Ir term in the central equation reverses too. The terminal voltage of the source being charged is then:

$$\mathbf{V = E + Ir} \quad (\textbf{during charging})$$

The terminal voltage now sits above the EMF, by exactly the Ir drop across the internal resistance. This is the only way the terminal voltage of a real source can exceed its EMF; in any other operating regime, the terminal voltage is always less than or equal to E.

Measuring EMF and internal resistance

To measure EMF and internal resistance accurately, one piece of instrument vocabulary must come first, which is input impedance. **The input impedance** of a voltmeter *is the resistance that the voltmeter itself presents to the circuit it is measuring. (*In **direct-current** work the word 'impedance' means the **same** thing as 'resistance'; we use the more general word to be consistent with the labels on commercial instruments, which often have alternating-current ranges where impedance and resistance differ.) A voltmeter's input impedance is important because, when the voltmeter is connected across a circuit element to measure its voltage, the voltmeter draws a small current of its own through its own input impedance, and that drawn current can disturb the reading.

A voltmeter is called high-impedance when its input impedance is so large compared with the resistances in the circuit being measured that the current it draws is negligible. For laboratory work, 'high' conventionally means at least a megaohm (a million ohms), and modern digital multimeters typically have input impedances of ten megaohm on their voltage ranges. A high-impedance voltmeter is the right instrument for measuring open-circuit EMFs because it disturbs the circuit it measures so little.

With that vocabulary in place, the cleanest way to measure the EMF of a source is the one already mentioned: connect a high-impedance (resistance) voltmeter across the source's terminals with no load attached. With $I = 0$ the relation $V = E - Ir$ collapses to $V = E$, and the voltmeter reads the EMF directly.

The residual current drawn by a ten-megaohm voltmeter from a 1.5V cell is $I = \frac{1.5V}{1010^6\Omega} = 1.5 \times 10^{-7}A$, which produces an Ir correction so small (a microvolt or so) that it cannot be detected on the voltmeter's own display. However, a more accurate method is the **potentiometer**, which we shall meet in a later section, and which draws no current at all at balance.

Measuring r requires a method known as the two-load method, which extracts both E and r from two pairs of (current, terminal-voltage) measurements taken at two different external loads. The method is built on the fact that the equation $V = E - Ir$ contains two unknowns (E and r), so we need two independent equations to solve for them, and we obtain those two equations by taking two measurements at two different loads.

Connect the source through a known external resistance R_1 and read the current I_1 with an ammeter. The closed-circuit equation around the loop (EMF equals sum of IR drops) gives:

$$E = I_1(R_1 + r)$$

Now change the external resistance to a different value R_2 and read the new current I_2:

$$E = I_2(R_2 + r)$$

These are two equations in two unknowns (E and r). Solve them simultaneously by subtracting one from the other to eliminate E:

$$0 = I_2R_2 + I_2r - I_1R_1 - I_1r$$
$$r(I_1 - I_2) = I_2R_2 - I_1R_1$$

Hence:

$$\mathbf{r = \frac{I_2R_2 - I_1R_1}{(I_1 - I_2)}}$$

Substituting back into either of the originals then gives E.

Now we turn to worked examples.

BINDER Example 17

A battery of EMF 12V and internal resistance 0.5Ω drives a current of 2A through an external circuit. Find the terminal voltage of the battery while the current is flowing.

Solution

The terminal voltage of a real source carrying a current is given by V = E − Ir, with the minus sign because the current is flowing in the natural direction (out of the positive terminal).

$$V = E - Ir$$

Substituting the given values:

$$V = 12V - (2A \times 0.5\Omega) = 11V$$

Making Sense of the Answer: *The terminal voltage is $11V$, which is $1V$ less than the EMF. That missing volt is the price the battery pays for its own internal resistance: every coulomb of charge loses one joule of energy on the way out of the cell, which goes into I^2r heat inside the casing. The external circuit therefore sees $11V$, not $12V$, and any device powered by this battery will operate as though it were connected to an ideal $11V$ source.*

Think Like a Physicist: *Always carry units through the substitution. Volts minus amperes-times-ohms gives volts (because $A \times \Omega = V$ is what Ohm's law says), and the unit check is a free verification that the formula has been applied correctly.*

HOT Example 18

A voltmeter is connected in parallel with a variable resistor R, which is in series with an ammeter and a cell. For one value of R the meters read 0.3A and 0.9V. For another value of R the meters read 0.25A and 1.0V. Find the value of R in each case, the EMF of the cell, and the internal resistance of the cell. State the assumption made about the resistance of the meters in the calculation.

Solution

Begin with the assumption.

We assume that the ammeter has negligible resistance compared with R and r (so its presence in the loop does not affect the current that the cell drives), and that the voltmeter has very high impedance compared with R (so the current it draws is negligible compared with the current through R, and the voltmeter reads the same as if it were not there).

Under these assumptions, the voltmeter reads the voltage across R, which is also the terminal voltage of the cell (since the only other thing in the loop is the ammeter, whose voltage drop is negligible).

The terminal-voltage equation $\mathrm{V} = \mathrm{E} - \mathrm{Ir}$ applies to both meter readings.
For the first reading:

$$0.9\mathrm{V} = \mathrm{E} - (0.3\mathrm{A})\mathrm{r}$$

For the second reading:

$$1.0\mathrm{V} = \mathrm{E} - (0.25\mathrm{A})\mathrm{r}$$

Subtract the first equation from the second:

$$1.0\mathrm{V} - 0.9\mathrm{V} = (0.3\mathrm{A} - 0.25\mathrm{A})\mathrm{r}$$
$$\mathrm{r} = 2\Omega$$

Substitute $\mathrm{r} = 2\Omega$ back into the first equation:

$$0.9\mathrm{V} = \mathrm{E} - (0.3\mathrm{A})(2\Omega)$$
$$\mathrm{E} = 0.9\mathrm{V} + 0.6\mathrm{V} = 1.5\mathrm{V}$$

The two external resistances follow from Ohm's law applied to R alone (each meter's reading on the voltmeter is the voltage across R, since the ammeter drop is negligible):

$$\mathrm{R}_1 = \frac{0.9\mathrm{V}}{0.3\mathrm{A}} = 3\Omega$$
$$\mathrm{R}_2 = \frac{1.0\mathrm{V}}{0.25\mathrm{A}} = 4\Omega$$

So the mystery cell has $\mathrm{E} = 1.5\mathrm{V}$ and $\mathrm{r} = 2\Omega$, and the two external resistances are 3Ω and 4Ω.

Making Sense of the Answer: *All four numbers are physically sensible. $E = 1.5V$ matches a standard zinc-carbon dry cell. The internal resistance $r = 2\Omega$ is rather large, indicating either an old cell or a cheap one; the cell loses a substantial fraction of its EMF to internal dissipation as soon as any current is drawn. At $I = 0.3A$ the cell loses $0.6V$ inside itself and delivers only $0.9V$ to the external circuit, an efficiency of just sixty percent. A higher-quality cell with $r = 0.1\Omega$ would deliver $1.47V$ at the same current, an efficiency of ninety-eight percent. Battery quality and internal resistance are inversely related.*

Think Like a Physicist: *Whenever you have two measurements and two unknowns, the two-equation method is almost always the right approach. Write down the same physical equation for each measurement, subtract one from the other to eliminate one unknown, solve for the other, then substitute back. The strategy is mechanical, but it requires you to identify in advance which physical equation is the same for both measurements. Here that equation was $V = E - Ir$, and the two unknowns E and r are properties of the cell that are the same in both readings; only V and I change.*

HOT Example 19

A 9V battery is tested by a student. With no load, a high-impedance voltmeter reads the terminal voltage as 9.0V. The battery is then connected to an external resistor; the voltmeter reading drops to 7.5V while the ammeter reads 0.5A. Find (i) the internal resistance of the battery; (ii) the external resistance; (iii) comment on what this experiment reveals about the quality of the battery.

Solution

(i) With no load, means that the circuit is an open circuit.

The open-circuit reading is $9.0\mathrm{V}$ with $\mathrm{I} = 0$, so the EMF is $\mathrm{E} = 9.0\mathrm{V}$.

With the external load connected, $V = E - Ir$ gives r:

$$7.5V = 9.0V - (0.5A)r$$
$$r = 3.0\Omega$$

(ii) Under load the voltmeter reads the terminal voltage, which is the same as the voltage across the external resistor R. Apply Ohm's law to R alone:

$$R = \frac{V}{I} = \frac{7.5V}{0.5A} = 15\Omega$$

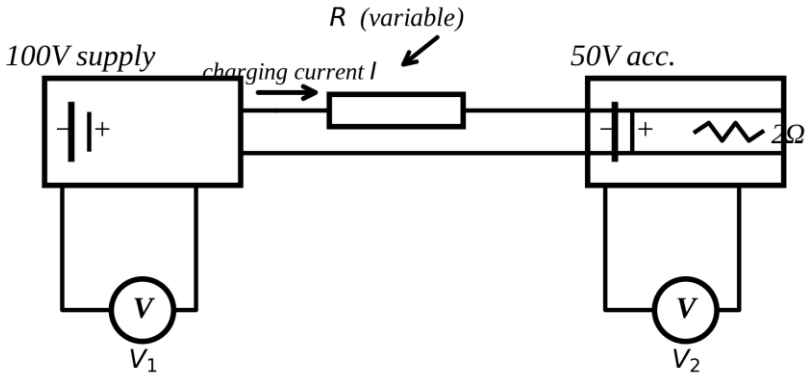
(iii) The experiment shows that the battery has a relatively high internal resistance, since its terminal voltage drops noticeably under load, meaning a significant amount of energy is lost inside the battery and it is not a very efficient or high-quality source for maintaining a steady voltage.

Making Sense of the Answer: *The internal resistance is $r = 3.0\Omega$, which is high for a $9V$ cell. Under the load tested, the battery delivers only $7.5V$ of its $9.0V$ EMF to the external circuit, which is $\frac{7.5}{9.0}$, or 83 percent efficiency. The remaining 17 percent of the available energy is dissipated inside the cell as heat. Battery quality scales inversely with r: the smaller r, the more of the EMF reaches the actual circuit.*

Think Like a Physicist: *The same battery, in different applications, can be 'good' or 'poor' depending on the load. A $9V$ smoke-detector battery with $r = 3\Omega$ might last years in service because a smoke detector draws only microamperes, at which current the Ir drop is negligible and the battery appears nearly ideal. The same battery in a $0.5A$ load shows its true weakness. Always evaluate r against the actual current drawn by the intended application; the number itself is meaningful only against a load context.*

REAL Example 20

An accumulator of EMF $50V$ and internal resistance 2Ω is being charged from a $100V$ DC supply, with a variable resistor R in series in the charging line. Two high-impedance voltmeters are connected, one across the supply terminals and one across the accumulator's terminals (see the figure below). Answer the following: (i) state the direction of the charging current; (ii) with $R = 0$, calculate the charging current; (iii) calculate the terminal voltage of the accumulator under that condition; (iv) predict how each voltmeter reading changes as R is increased from zero.



Solution

(i) The supply is the more powerful source. It must force current into the positive terminal of the accumulator, against the direction the accumulator would drive on its own. The charging current therefore flows from the positive terminal of the supply, through the variable resistor R, into the positive terminal of the accumulator, out of the accumulator's negative terminal, and back to the negative terminal of the supply.

(ii) The cleanest way into this problem is to think of each component's terminal voltage separately. Three voltages are in play.

First, the supply. The phrase '100V DC supply' is the standard shorthand for an idealised voltage source: a source whose terminal voltage is held steady at 100V regardless of the current it delivers (equivalently, a source with negligible internal resistance). We adopt that idealisation throughout this problem.

Second, the accumulator. Being charged, its terminal voltage obeys the charging formula derived earlier in this section: $V = E + Ir = 50V + I(2\Omega) = 50V + (2\Omega)I$.

Third, the connecting wires. With $R = 0$, there is nothing else in the loop. The supply's positive terminal is connected directly to the accumulator's positive terminal by wires of negligible resistance, and similarly for the negative terminals. Two points joined by a wire of zero resistance must be at the same potential (because Ohm's law applied to that wire reads $V = I \times 0 = 0$ for any current). So the supply's terminal voltage equals the accumulator's terminal voltage:

$$V_{supply} = V_{acc}$$
$$100V = 50V + (2\Omega)I$$
$$I = \frac{100V - 50V}{2\Omega} = 25A$$

(iii) Substituting $I = 25A$ into the charging formula gives the accumulator's terminal voltage:

$$V = E + Ir = 50V + (25A)(2\Omega) = 100V$$

The accumulator's terminal voltage matches the supply voltage exactly, as we already noted in part (ii). With $R = 0$ there is nothing else in the loop, so the supply's $100V$ must appear directly across the accumulator's terminals.

(iv) When R is no longer zero, R sits in series between the supply and the accumulator. Ohm's law gives the voltage drop across R as $V_R = IR$. The supply's $100V$ is then shared (in series) between R and the accumulator's terminals:

$$V_{supply} = V_R + V_{acc}$$
$$100V = IR + (50V + (2\Omega)I)$$

Rearranging for I:

$$I = \frac{50V}{R + 2\Omega}$$

which falls as R rises. The supply voltmeter V_1 reads $100V$ regardless of R; the supply's own terminal voltage is independent of R (the supply is assumed to be a perfect $100V$ source). The accumulator voltmeter V_2 reads $V = E + Ir = 50V + (2\Omega)I$. As I falls toward zero with rising R, V_2 falls toward $E = 50V$. In the limit of very large R, almost no charging current flows, and the accumulator voltmeter reads its EMF directly.

Making Sense of the Answer: *The charging current at $R = 0$ is $25A$, which is dangerously large. No real charging room would allow this. Practical battery chargers always include a series resistance or, more commonly, an active current-limiting circuit that holds the charging current to a safe value (typically a few amperes for a small lead-acid cell, up to twenty or thirty for a large traction battery). The whole point of the variable R in this example is to let the operator dial the charging rate to a safe number. Increasing R slows the charge; the engineering trade-off is between charging quickly and the heat the cells can safely shed.*

Think Like a Physicist: *The cleanest way through any single-loop circuit containing two EMF sources is to think in terms of each component's terminal voltage, not in terms of energy balance. The supply's terminal voltage is a known property of the supply (here, just $100V$ since the supply is ideal). The accumulator's terminal voltage obeys $V = E + Ir$ during charging or $V = E - Ir$ during discharge. The voltage drop across any resistor is IR. The supply's voltage is then distributed across the loop's series elements, and the distributed voltages add up to the supply's voltage. This gives one equation in one unknown (the current), and the algebra is trivial.*

The Trap to Avoid: *The common error in charging problems is to forget that the accumulator's terminal voltage exceeds its EMF (it is $E + Ir$, not $E - Ir$). To decide which sign to use, ask: in which direction is the current passing through this source? In the direction the source would drive on its own (i.e. out of its positive terminal), use $V = E - Ir$. Forced backwards through the source (into its positive terminal, i.e. being charged), use $V = E + Ir$. Get the sign wrong and a $25A$ answer turns into a $75A$ answer, a factor of three off.*

HOT Example 21

A careless mechanic at a Dar es Salaam garage is leaning over the open bonnet of a car, working on the engine. He drops a metal spanner from his hand, and the spanner falls across the two terminals of the 12V car battery and rests there, with its body bridging the positive and negative terminals directly. The battery has internal resistance r = 0.04Ω. (a) Compute the current that flows through the spanner. (b) Compute the power dissipated inside the battery at that current. (c) Predict what happens physically in the next few seconds, and explain why a car battery has been designed in a way that makes this situation particularly dangerous.

Solution

(a) The spanner is a piece of metal of negligible resistance, so it presents essentially zero external resistance between the battery's two terminals. The voltage drop across the spanner is therefore zero (by Ohm's law applied to the spanner: $V = I \times 0 = 0$), and since the spanner is the only thing connected between the battery's terminals, the terminal voltage of the battery is $V = 0$. Apply $V = E - Ir$:

$$0 = 12V - I(0.04\Omega)$$
$$I = \frac{12V}{0.04\Omega} = 300A$$

(b) Power dissipated inside the battery is $P = I^2r$:

$$P = (300A)^2(0.04\Omega) = 3600W$$

(Equivalently, the total power delivered by the battery is $P = EI = (12V)(300A) = 3600W$, all of which is dissipated internally because the external resistance is zero.)

(c) The current of 300A obtained in part (a) is what is called the **short-circuit current** of the battery: the current that flows when the external resistance is approximately zero. The spanner has acted as the short-circuiting wire.

The 3600W of power is concentrated inside the small volume of the battery's interior. In the first fraction of a second this heats the electrolyte rapidly. The point of contact between the spanner and the battery terminals carries the full 300A through a tiny surface area; the local resistance there (small but not exactly zero) dissipates further power, and the spanner glows red hot and can melt and weld itself to the terminal. The hydrogen and oxygen produced inside the casing by the rapid electrolysis can ignite. In the worst case the battery's casing ruptures and sprays sulphuric acid.

The reason this happens to a car battery and not, so spectacularly, to a torch cell is the design choice that gives the car battery its usefulness in the first place. A car's starter motor needs to draw hundreds of amperes for a couple of seconds while the engine is cranked. The battery must therefore have very low internal resistance, so that the Ir drop inside the battery does not collapse the terminal voltage when those high currents are drawn. That same low r is what permits the catastrophic 300A short-circuit current. Useful and dangerous are two sides of the same design.

Making Sense of the Answer: *The short-circuit current scales inversely with r. A car battery ($E = 12V$, $r = 0.04\Omega$) gives $300A$. A standard zinc-carbon dry cell ($E = 1.5V$, $r = 0.5\Omega$) gives only $3A$. The power dissipated scales as $\frac{E^2}{r}$, so the car battery dissipates $3600W$ in a short, while the dry cell dissipates only $4.5W$. That is a factor of 800 difference in danger, even though the EMFs differ by a factor of only 8. Internal resistance dominates the danger calculation.*

Think Like a Physicist: *The short-circuit current is a property of the source. It tells you the stiffness of the source: a source with small r holds its terminal voltage steady under load (because the Ir correction is small), and equivalently has a large short-circuit current. The two are inverse statements of the same property. Whenever a manufacturer quotes the internal resistance of a battery, they are giving you, implicitly, the short-circuit current, and through it the maximum power the battery could ever deliver to an external load.*

Five examples behind us.

Every example so far has had a single source driving a single resistor. Real circuits are not so simple. A torch has two cells stacked in series; a smartphone has parallel cells inside its casing; a household lighting circuit has dozens of lamps in parallel; and an oscillator on a teaching bench may have a Wheatstone bridge made of four resistors arranged with mirror symmetry. The next section opens that catalogue. Series resistors, parallel resistors, series cells, parallel cells, the symmetry methods that simplify complicated networks, and the shunt-and-multiplier modifications that turn a basic galvanometer into a useful ammeter or voltmeter, all follow. The same equations $V = IR$ and $V = E - Ir$ hold throughout; what changes is the bookkeeping.

RESISTORS, CELLS, SYMMETRY, AND METER MODIFICATION

Every circuit so far in this chapter has had a single source driving a single resistor. Real circuits are not so simple. The torch in your pocket has four cells stacked end to end. The phone in your other pocket has parallel cells inside its casing. The lighting circuit of your house has a dozen lamps in parallel between the live and neutral rails. The oscillator on your physics teacher's bench may have a Wheatstone bridge made of four resistors with a fifth one hung across its middle. None of these can be solved by the single-loop method of the last section.

The good news is that the underlying physics has not changed. Ohm's law still gives the voltage across each resistor, the EMF-plus-internal-resistance relation still gives the behaviour of each cell, and the conservation laws (charge in equals charge out at every junction, energy in equals energy out around every loop) still hold. What changes is the bookkeeping. This section catalogues the cases that yield to simple bookkeeping: resistors in series, resistors in parallel, cells in series, cells in parallel, networks that reduce by their geometric symmetry, and the meter modifications that turn a basic galvanometer into a useful ammeter or voltmeter. The next section, on Kirchhoff's laws, picks up everything that this one cannot reduce.

Series resistors

Three resistors of resistances R_1, R_2, and R_3 are connected end to end in a single closed loop, with a source of voltage V driving a current I round the loop. The three resistors are said to be in series. The following figure shows the configuration.

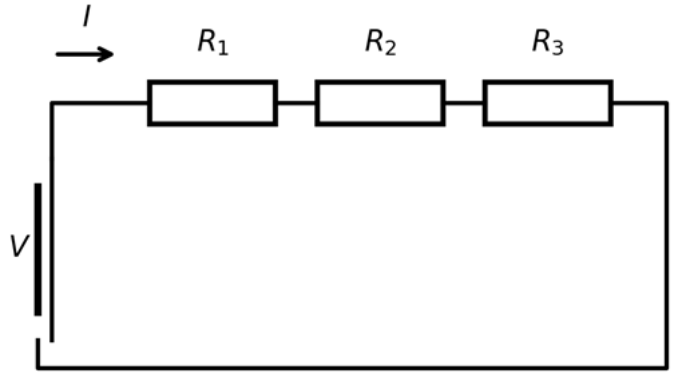


Figure: ***Three resistors connected in series with a source.*** *The same current I passes through each resistor. The total voltage V applied across the series combination equals the sum of the voltage drops across the individual resistors.*

Two physical facts about this arrangement are worth saying carefully, because everything else follows from them. The first is that the current is the same through every resistor. Charge cannot accumulate at any point in the loop (where would it go?), so the rate at which charge flows past any cross-section of the loop is the same as the rate at which it flows past any other cross-section. The second fact is that the total voltage drop around the loop equals the source voltage. This is conservation of energy: the work done per coulomb by the source must equal the energy lost per coulomb in the resistors, because the resistors are the only places the energy can be lost.

Apply Ohm's law to each resistor in turn, using the same current I because the current is the same everywhere:

$$V_1 = IR_1, \quad V_2 = IR_2, \quad V_3 = IR_3$$

Add the three equations together. The left-hand sides sum to the total voltage V (by conservation of energy); the right-hand sides factor I out:

$$V = V_1 + V_2 + V_3 = I(R_1 + R_2 + R_3)$$

Compare this with the Ohm's-law statement $V = IR_{eq}$ that would apply to a single equivalent resistor R_{eq} carrying the same total current I under the same source voltage V.

$$IR_{eq} = I(R_1 + R_2 + R_3)$$

$$R_{eq} = R_1 + R_2 + R_3$$

The same argument works for any number of resistors in series, not just three.
For n resistors R_1, R_2, ..., R_n in series:

$$\mathbf{R_{eq} = R_1 + R_2 + \cdots + R_n}$$

There is a useful consequence about how the source voltage is shared between the resistors. The voltage across each resistor is $V_i = IR_i$, and the current I is the same in every one. So V_i is directly proportional to R_i: the largest resistor gets the largest voltage drop, the smallest resistor gets the smallest.

Parallel resistors

Three resistors of resistances R_1, R_2, and R_3 are connected in parallel between two common points (call them the top rail and the bottom rail). A source of voltage V is connected across the rails. The following figure shows the configuration.

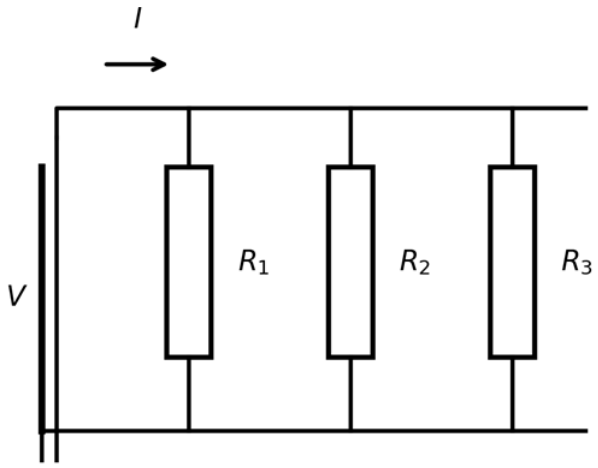


Figure: ***Three resistors connected in parallel with a source.*** *The same voltage V appears across each resistor. The total current I drawn from the source equals the sum of the currents through the individual resistors.*

The two physical facts to keep in mind are now the opposite of the series case. The voltage across each resistor is the same (because each resistor is connected between the same two rails, and the potential difference between two points cannot depend on which path you follow). The total current drawn from the source is the sum of the currents through the individual resistors (because at the junction where the rails split into the three branches, charge in equals charge out).

Apply Ohm's law to each resistor with the common voltage V:

$$I_1 = \frac{V}{R_1}, \quad I_2 = \frac{V}{R_2}, \quad I_3 = \frac{V}{R_3}$$

Add the three currents to give the total drawn from the source:

$$I = I_1 + I_2 + I_3 = V\left(\frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3}\right)$$

Compare this with Ohm's law applied to a single equivalent resistor: $I = \frac{V}{R_{eq}}$. The two read identically when:

$$\frac{1}{R_{eq}} = \frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3}$$

The same argument generalises to n parallel resistors in the obvious way.

$$\mathbf{\frac{1}{R_{eq}} = \frac{1}{R_1} + \frac{1}{R_2} + \frac{1}{R_3} + \cdots + \frac{1}{R_n}}$$

The special case of two resistors in parallel is so common that the formula is worth remembering in its compact form. Setting $\frac{1}{R_{eq}} = \frac{1}{R_1} + \frac{1}{R_2}$ and rearranging gives:

$$R_{eq} = \frac{R_1 R_2}{R_1 + R_2}$$

The current through each branch is inversely proportional to its resistance. The smallest resistor in a parallel combination takes the largest share of the current; the largest resistor takes the smallest share. For n equal resistors connected in parallel, the current is shared equally among them, and the equivalent resistance is equal to the resistance of one resistor divided by n, as shown below:

$$\frac{1}{R_{eq}} = \frac{1}{R} + \frac{1}{R} + \frac{1}{R} + \cdots = \frac{1 + 1 + 1 + \cdots}{R} = \frac{n}{R}$$

$$\mathbf{R_{eq} = \frac{R}{n}}$$

Series cells

Cells in series follow the same logic as resistors in series, with one extra fact: each cell adds its own EMF, with sign, to the total EMF of the chain. Suppose n identical cells, each of EMF E and internal resistance r, are connected in series with all the positive terminals joining to the negative terminals of their neighbours (so that the EMFs all push in the same direction round the loop). The figure below shows the arrangement.

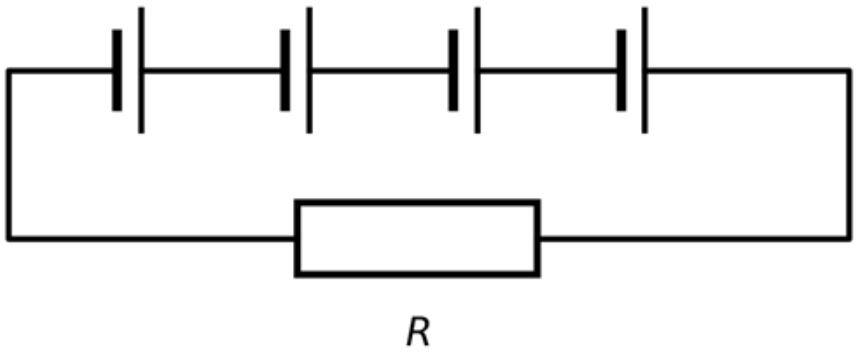


Figure: n identical cells in series, all driving in the same direction*. The total EMF is nE and the total internal resistance is nr.*

Because the cells are in series, the same current passes through every cell. Each cell's internal resistance r drops a voltage Ir along the chain, and each cell's EMF E lifts the potential by E. The two effects add cell by cell, giving:

$$\mathbf{E_{eq} = nE, \quad r_{eq} = nr}$$

The equivalent EMF is n times the EMF of a single cell; the equivalent internal resistance is n times the internal resistance of a single cell. Both scale with the number of cells. Series stacking is the right choice when you want a higher voltage than a single cell can supply, and you can afford the larger internal resistance that comes with it. The 6V torch in your house is a stack of four 1.5V cells in series; the 12V battery in a small motor-vehicle is six 2V lead-acid cells in series. The 1.5V is the

natural EMF of a zinc-carbon cell and the 2V is the natural EMF of a lead-acid cell; cells are stacked to multiply the voltage by a small integer until the practical voltage required is reached.

Parallel cells of equal EMF

Cells in parallel work in the opposite direction. Suppose n identical cells (same EMF E, same internal resistance r) are connected in parallel, all positive terminals joining to a common top rail and all negative terminals joining to a common bottom rail. The next figure shows the arrangement.

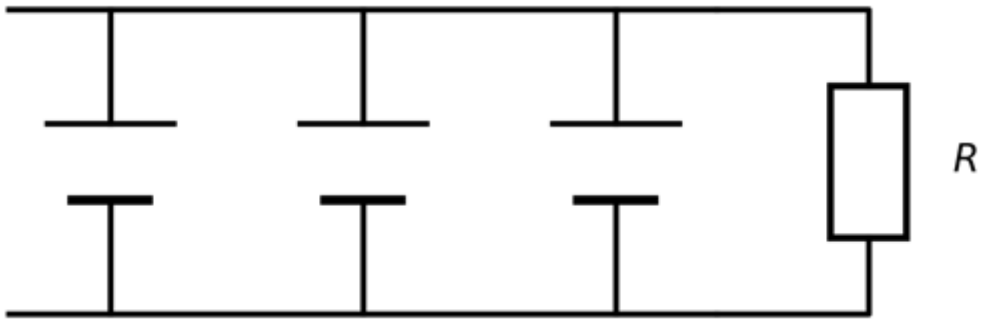


Figure: ***n identical cells in parallel.*** *The equivalent EMF is unchanged at E; the equivalent internal resistance is r divided by n. Voltage is unchanged, but the current capability is multiplied by n.*

Because all n cells share the same two rails, they all have the same terminal voltage; and because they all have the same EMF E, they all drive in the same direction. The equivalent EMF of the combination is therefore just E itself, the EMF of any single cell. The internal resistances, however, combine as resistors in parallel: n equal resistances r in parallel give $\frac{r}{n}$. So:

$$E_{eq} = E, \quad r_{eq} = \frac{r}{n}$$

The voltage stays the same, but the internal resistance falls by a factor of n. This is useful when you want to draw more current than a single cell can comfortably supply, at the cell's natural voltage. A **smartphone battery** pack is several lithium cells of identical EMF wired in parallel; the voltage stays at the cell's natural value (about 3.7V) while the current capability multiplies up.

The two strategies (series for voltage, parallel for current) are often combined: a large vehicle traction battery is a series-parallel grid of many cells together delivering both the high voltage and the high current demands of the motor.

If the parallel cells do not have equal EMFs, the analysis is less elementary. The cells then drive current internally from one to the other even with no external load, wasting energy. The general parallel-cell formula is derived in the next section using Kirchhoff's laws, and the practical advice is to avoid the configuration unless you specifically intend it.

Symmetry methods in networks

The series and parallel formulas of the last four subsections handle every network that can be broken into nested chunks of pure series or pure parallel structure. Many real circuits look exactly like this, and they reduce by repeated application of the two formulas alone. But not every circuit is so cooperative.

Consider a cubical lattice of twelve identical resistors, one on each edge of the cube, and connect a battery across the cube's main diagonal. *What is the equivalent resistance between the two diagonally opposite corners?* Stare at it as long as you like, you will not find any pair of edges that is in pure series with another, nor any group of edges in pure parallel with the rest. The series and parallel formulas alone are stuck.

There is, however, a way through. The cube has hidden order: it looks the same from several different points of view. That geometric symmetry leads, through a short argument we shall present next, to the conclusion that certain corners of the cube must be at the same potential. Once we know that, we may join those corners with a wire (because the wire carries no current) and redraw the network. The redrawn network turns out to be a simple series chain of parallel blocks, which the series and parallel formulas of the earlier subsections reduce in one line. That is what symmetry methods are. The next few paragraphs spell out the logic step by step.

Why symmetry forces two nodes to be at the same potential

Suppose you can perform some operation on the network (a rotation, a reflection, or a relabelling of corners) that has two properties. First, after the operation the picture looks identical to what it was before. Same wires, same resistors, same battery in the same place. Second, the operation swaps two specific nodes, call them X and Y. Then the operation has the unavoidable consequence that $V_X = V_Y$.

The argument is short. Before the operation, node X sits at some potential V_X and node Y sits at some other potential V_Y. After the operation the network is physically the same network as before (that is what the first property tells us), so the potential at

any chosen point of it must be the same as before. But the operation has swapped X and Y, so the position now called 'X' used to be called 'Y'. That position must therefore have both potentials at the same time: the one it had before the operation (V_Y, from its old name) and the one it has after the operation (V_X, from its new name, since the picture is unchanged). The only way both numbers can describe the same point is if they are equal.

In plain words: *if a circuit looks the same after we swap two of its nodes, the two nodes cannot tell each other apart electrically, so they must be at the same potential.* This is the central lemma of symmetry methods.

Why we may merge equipotential nodes

Once we know that $V_X = V_Y$, we are allowed to add a wire of zero resistance between X and Y. The added wire does not change any current anywhere in the network. Why not? Because the current that would flow in the added wire would obey Ohm's law applied to the wire alone: $I = V/R$. The voltage V across the wire is the difference $V_X - V_Y$, which we have just shown is zero. The resistance R of the added wire is also zero (it is an ideal wire). So the current is $I = 0/0$ (indeterminate from this calculation alone), but the rest of the circuit has not changed at all, so the current actually carried by the added wire is whatever amount the rest of the network needed to be carried between X and Y anyway, namely none. Adding the wire is electrically the same as fusing the two nodes into one combined node. We say the two equipotential nodes have been merged.

Stated as a rule: *nodes proved equipotential by symmetry can be merged into a single node without changing anything that matters.*

A warm-up example: the balanced Wheatstone bridge

Before we attack the cube, here is a smaller example that uses the same two steps. The figure below shows four resistors of equal resistance R arranged as the four sides of a diamond ABCD, with a galvanometer G connected across the diagonal BD. A battery drives current in at A and out at C.

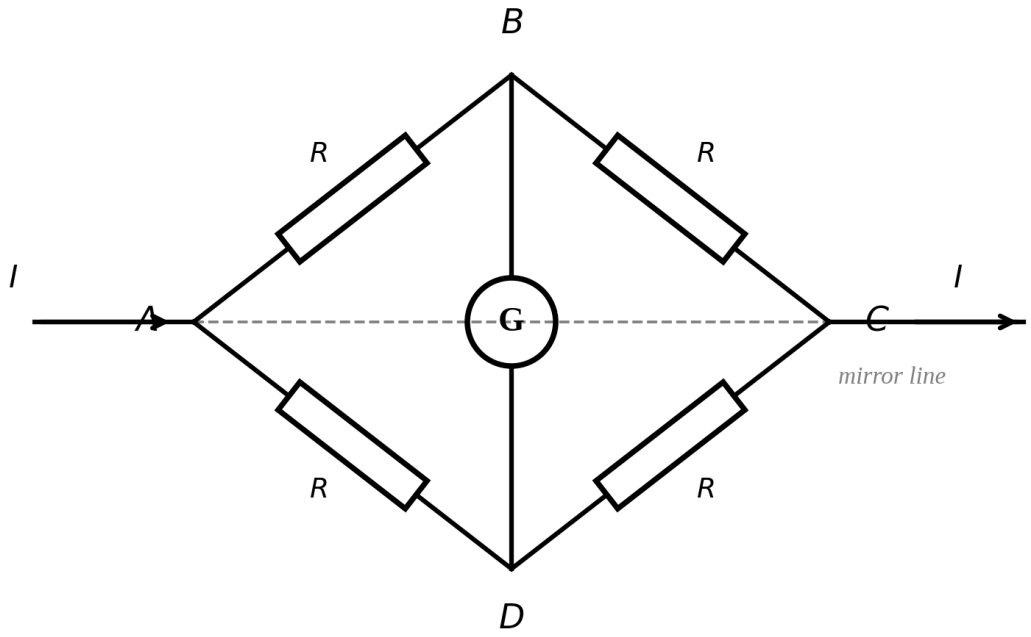


Figure: ***A balanced Wheatstone bridge****: four equal resistors R on the sides of a diamond, with a galvanometer G across the diagonal BD. The input is at A and the output is at C. The dashed line through A and C is the mirror axis of the network.*

Find a symmetry of this network. Reflect the whole figure across the horizontal dashed line through A and C. The top half (the path A→B→C through two R resistors) flips down to where the bottom half used to be (the path A→D→C through two R resistors). The bottom half flips up. After the reflection the picture looks identical to what it was before: the resistors are still all R, the battery is still at A and C, and the galvanometer is still on the vertical line. The reflection has swapped B with D, while leaving the picture unchanged.

The lemma above then tells us that B and D must be at the same potential: $V_B = V_D$. The merge rule above then tells us we are allowed to merge B and D into a single node. The galvanometer G, which sits between B and D, therefore has zero voltage across it and carries zero current. The galvanometer reads zero; this is the classic balanced-bridge condition that turns the Wheatstone bridge into a precision instrument for comparing resistances.

With B and D merged into one node, the four resistors are no longer a bridge but a much simpler combination: AB and AD are now in parallel between A and the merged BD node, and BC and DC are in parallel between the merged node and C. Each

parallel pair has equivalent resistance $R/2$, and the two pairs are then in series along the path A to merged BD to C, giving the total resistance between A and C as:

$$\frac{R}{2}+\frac{R}{2}=R$$

Notice how the two-step recipe (find a symmetry, merge the swapped nodes) reduced the bridge to a straightforward series-parallel calculation. The same recipe, applied to a more elaborate network, is what we use on the cube.

The cube of twelve resistors

Now back to the cube. The cube has eight corners; label them A (the input), B, D, E (the three corners directly adjacent to A, joined to it by single edges), C, F, H (the three corners directly adjacent to G, joined to it by single edges), and G (the output, diagonally opposite A across the cube). Each of the twelve edges is a resistor of resistance R.

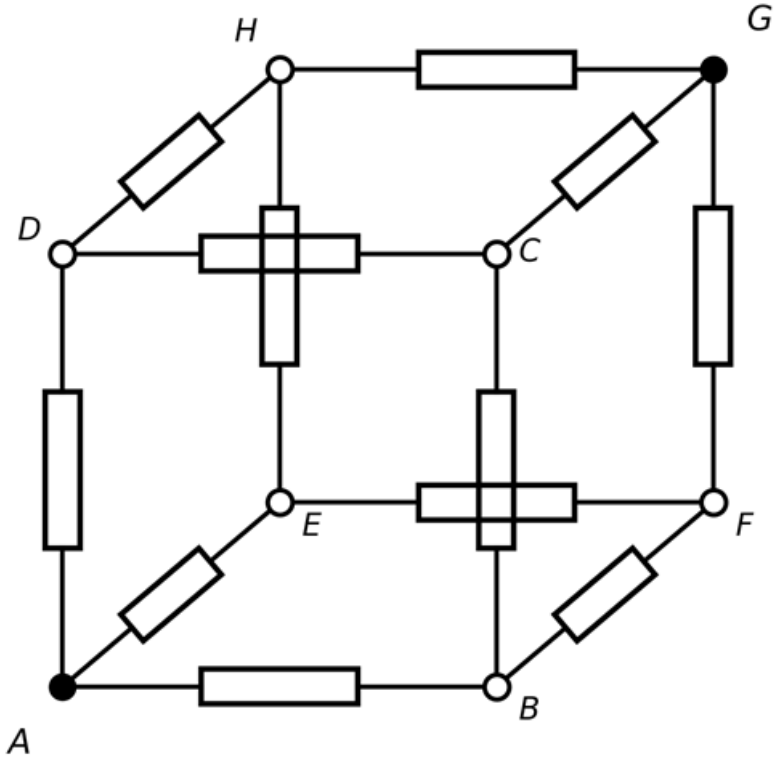


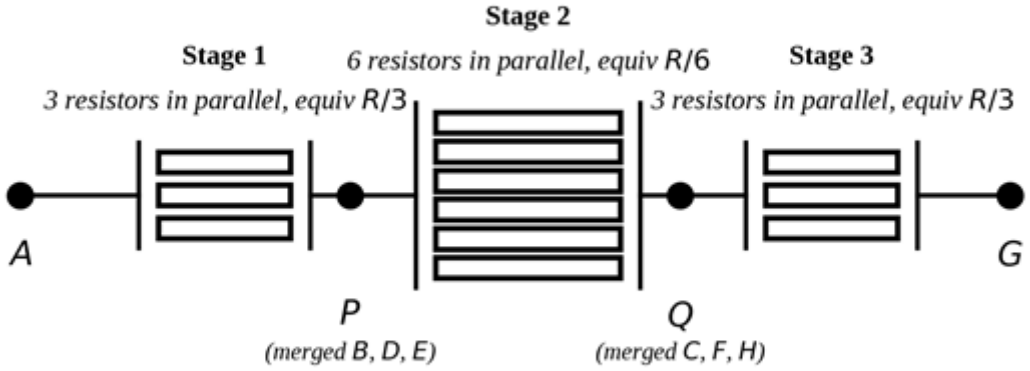
Figure: The cubical network of twelve equal resistors. The eight corners are labelled A, B, C, D, E, F, G, H. Current enters at corner A (filled circle, bottom-left) and leaves at the diagonally opposite corner G (filled circle, top-right). The three corners B, D, E adjacent to A and the three corners C, F, H adjacent to G are drawn as open circles. Each of the twelve edges is a resistor of resistance R.

Find the symmetry. Pick the cube up by the two corners A and G (the diagonal along which current enters and leaves) and rotate it through 120° about that diagonal. The three corners B, D, E cycle among themselves: B moves to where D was, D moves to where E was, E moves to where B was. The three corners C, F, H cycle similarly. A stays at A, G stays at G, every resistor moves to the position previously occupied by another resistor (with the same resistance R, since they are all identical). So after the rotation, the picture looks identical to what it was before.

Apply the lemma above. The rotation cycled B to D, D to E, E to B. So all three of B, D, E must sit at the same potential. By the same argument applied on the output side, all three of C, F, H must also sit at a (common) potential of their own. We now have two groups of three equipotential nodes each.

Apply the merge rule above. Merge B, D, E into a single node, call it P. Merge C, F, H into a single node, call it Q. The cube has now collapsed into a much simpler picture, which the next figure draws explicitly.

After the symmetry merge, the cube becomes a series chain of three parallel blocks



$$R_{eq} = \frac{R}{3} + \frac{R}{6} + \frac{R}{3} = \frac{5R}{6}$$

Figure*: **The cube AFTER the symmetry merge.** The three edges from A to its neighbours now connect A to the single merged node P, so they are three resistors R in parallel (equivalent R/3). The six edges that connected an A-neighbour to a G-neighbour now connect P to Q, so they are six resistors R in parallel (equivalent R/6). The three edges from the G-neighbours to G now connect Q to G, so they are three resistors R in parallel (equivalent R/3). These three blocks lie in series along the path A→P→Q→G, and their resistances add: total = R/3 + R/6 + R/3 = 5R/6.*

The figure makes the conclusion obvious: the cube, after merging, is just three parallel blocks in series. The first block has three equal R resistors in parallel (equivalent R/3), the middle block has six equal R resistors in parallel (equivalent R/6), and the last block has three equal R resistors in parallel (equivalent R/3). The series chain adds:

$$\mathbf{R_{eq} = \frac{R}{3} + \frac{R}{6} + \frac{R}{3} = \frac{5R}{6}}$$

The cube's twelve-resistor network has been reduced to a single equivalent resistance of 5R/6. Example 22 below plugs in the numerical R = 2Ω and a 6V battery to find the total current drawn.

Summary recipe

Symmetry methods always follow the same four steps. Each step is set out below with its own short explanation of how to do it in practice.

Step 1: Find a symmetry of the network

Look for an operation you can perform on the picture that leaves it identical. The operation may be a rotation, a reflection across a line, or a permutation of node labels. The test of success is this: after the operation, every wire still goes between the same two corners as before, every resistor still has the same resistance as before, and the battery still sits at the same two terminals as before. The picture must look identical including the battery, not only the resistors.

In the balanced-bridge example above, the symmetry was the reflection across the horizontal line through A and C. In the cube example above, the symmetry was a 120-degree rotation about the diagonal from A to G. These are the two most common kinds of symmetry you will meet at this level: a mirror reflection and a rotational permutation.

Step 2: Identify the nodes that the symmetry moves

Track each node of the network through the operation. Some nodes (typically the input and output of the network) will stay in place. Other nodes will be swapped, or cycled, with one or more of their neighbours. List the groups of nodes that get cycled among themselves by the operation. Each such group has been proved equipotential by the lemma proved earlier.

In the bridge, the reflection swapped B with D, so {B, D} formed one equipotential group. In the cube, the 120-degree rotation cycled B to D to E to B, and cycled C to F to H to C, so {B, D, E} and {C, F, H} formed two equipotential groups. The number of groups equals the number of independent merges you will perform in the next step.

Step 3: Merge each equipotential group into a single node

For each equipotential group identified in Step 2, draw a fresh node and connect every member of the group to it with a zero-resistance wire. By the merge argument given earlier, these added wires carry no current, so no measurable quantity has been changed. The old nodes can now be erased and replaced by the single new node, which inherits all the resistors that previously hung off any member of the group.

Give each merged node a short fresh name (P, Q, and so on). This keeps the redrawn picture clean and avoids the confusion of carrying three labels for one electrical point. In the cube, the merged A-neighbour node became P and the merged G-neighbour node became Q.

Step 4: Redraw the simplified network, then reduce by series and parallel

Sketch the simplified network with the merged nodes in place. Count how many of the original resistors now run between each pair of nodes; resistors that previously ran between two members of the same equipotential group will have collapsed to short wires and may be ignored, while resistors that ran from one group to another are all now in parallel between the corresponding merged nodes.

Replace each parallel cluster by its equivalent resistance using the parallel-resistor formula introduced earlier. The resulting picture is a simple series chain of equivalent resistors, and the series-resistor formula introduced earlier finishes the job. The total equivalent resistance comes out in one or two further lines of arithmetic.

In the cube, after the merge there were three resistors in parallel between A and P, six resistors in parallel between P and Q, and three resistors in parallel between Q and G. The three parallel clusters had equivalent resistances R/3, R/6, and R/3 respectively, and these three sat in series along the path from A to G, summing to 5R/6.

Two practical warnings

First, the symmetry must respect the battery as well as the resistors. An operation that moves the battery to a different position on the network does not count, because the picture is no longer identical. This is why the balanced bridge has a mirror axis through A and C (the battery terminals) and not through B and D: a reflection across the BD axis would carry the battery to a different pair of terminals, which is a different physical setup.

Second, not every circuit has a useful symmetry. Networks like the unbalanced Wheatstone bridge of worked Example 28 have no symmetry that swaps the bridge nodes, so this method does not apply to them. For those circuits we shall need the more powerful method of the next section, Kirchhoff's laws.

Meter modification: shunts and multipliers

Two of the most common measuring instruments in any physics laboratory are the ammeter and the voltmeter. Both are built up from the same basic component: a sensitive **galvanometer**, which *is a current-measuring instrument whose pointer deflects in proportion to the current passing through its coil.* A typical laboratory galvanometer gives full-scale deflection at a current of milliamperes, with the coil of resistance fifty to one hundred ohms. The galvanometer by itself is therefore neither an ammeter (which needs a range of amperes, not milliamperes) nor a voltmeter (which reads voltage, not current). To turn it into either, you add a single extra resistor in the right place.

To convert the galvanometer into an ammeter of full-scale range $\mathbf{I_{max}}$, connect a small resistor **S** (the **shunt**) in parallel with the galvanometer. The shunt diverts most of the current around the galvanometer, so that the galvanometer carries only its own small full-scale current $\mathbf{I_g}$ while the total current through the combination reaches I_{max}. At full-scale deflection, I_g passes through the galvanometer and $I_{max} - I_g$ passes through the shunt. The voltage across the galvanometer is I_gG (with G the galvanometer's own resistance), and this is the same as the voltage across the shunt (they are in parallel), so:

$$I_gG = (I_{max} - I_g)S$$

Solving for the required shunt:

$$\mathbf{S = \frac{I_gG}{I_{max} - I_g}}$$

Because I_g is small compared with I_{max}, the shunt resistance S comes out very small (a fraction of an ohm in typical cases). The full ammeter (galvanometer in parallel with shunt) has effective resistance equal to the parallel combination of G and S, which is dominated by the smaller of the two, namely S. This is exactly what an ammeter needs: *a very low resistance, so that placing the ammeter into a circuit hardly changes the current already flowing there. A good ammeter is a near-perfect conductor.*

To convert the same galvanometer into a voltmeter of full-scale range $\mathbf{V_{max}}$, connect a large resistor **M** (the **multiplier**) in series with the galvanometer. At full-scale deflection the current through the combination is $\mathbf{I_g}$ (the galvanometer's full-scale current), and the voltage across the combination is V_{max}. By Ohm's law applied to the (galvanometer + multiplier) series chain:

$$V_{max} = I_g(G + M)$$

Solving for the required multiplier:

$$\mathbf{M = \frac{V_{max}}{I_g} - G}$$

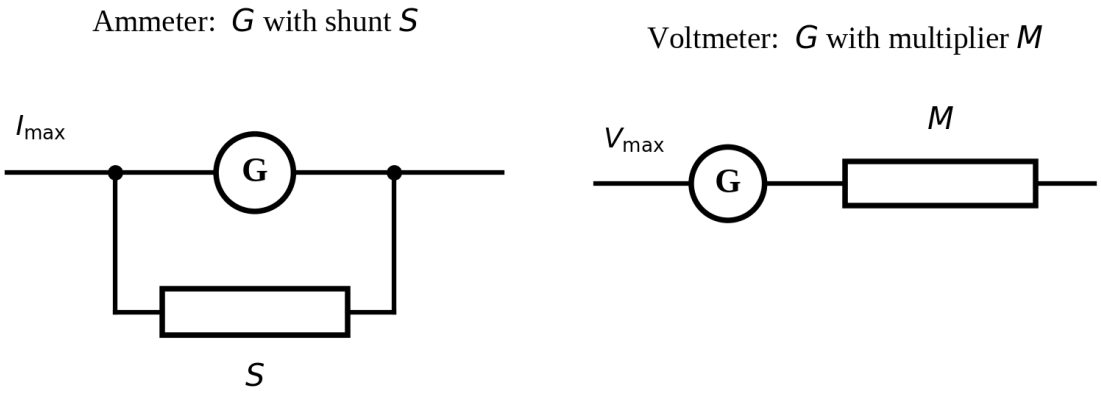


Figure: ***Galvanometer modifications.*** *Left: an ammeter, made by adding a small shunt S in parallel with the galvanometer G. Right: a voltmeter, made by adding a large multiplier M in series with G.*

Because V_{max} is typically several volts and I_g is only a milliampere, the multiplier resistance comes out large (thousands of ohms or tens of thousands of ohms). The full voltmeter has effective resistance $G + M$, which is essentially M. This is what a voltmeter needs: *a very high resistance, so that placing the voltmeter across a circuit element draws almost no current from the circuit and hardly changes the voltage already present there. A good voltmeter is a near-perfect insulator.* Worked Examples 23 and 24 design a shunt and a multiplier respectively for a typical laboratory galvanometer.

Seven examples follow. Each one puts a different method of this section to the test, and the last anticipates a tool we shall need in the next section.

BINDER Example 22

A cubical network is built from twelve resistors of 2Ω each, one on each of the twelve edges of the cube. The network is connected across its main diagonal (from corner A to the diagonally opposite corner G) to a 6V battery of negligible internal resistance. Find the equivalent resistance of the network and the total current drawn from the battery.

(The cubical network is the same as the one used in the symmetry-methods subsection, with $R = 2\Omega$ on each edge.)

Solution

The full derivation was carried out in the symmetry-methods subsection under the heading 'The cube of twelve resistors'. The three-fold rotational symmetry about the main diagonal $A \rightarrow G$ forces the three A-neighbours B, D, E to share a common potential, and likewise the three G-neighbours C, F, H. Merging each group into a single node and redrawing collapses the cube into three parallel blocks in series, with equivalent resistance:

$$R_{eq} = \frac{R}{3} + \frac{R}{6} + \frac{R}{3} = \frac{5R}{6}$$

With $R = 2\Omega$:

$$R_{eq} = \frac{5 \times 2\Omega}{6} = 1.67\Omega$$

The total current drawn from the 6V battery (negligible internal resistance, so terminal voltage equals EMF) follows from Ohm's law applied to the equivalent resistance:

$$I = \frac{V}{R_{eq}} = \frac{6V}{1.67\Omega} = 3.6A$$

Making Sense of the Answer: *Symmetry has done all the work. The twelve-resistor cube, which has no obvious series or parallel decomposition when stared at directly, collapses in three short stages once the symmetry merge is performed. The current* $3.6A$ *is consistent with the small total resistance: a* $6V$ *battery into less than two ohms must drive a few amperes.*

Think Like a Physicist: *Always look for symmetry first when a network does not reduce by obvious series and parallel. The four-step recipe given earlier is general: find an operation that leaves the picture unchanged, identify the nodes it swaps or cycles, merge them, redraw and reduce. When no symmetry can be found, the Kirchhoff's laws of the next section give a universal method, but they are also the slower one. Symmetry first, Kirchhoff only as a fallback.*

HOT Example 23

A galvanometer has resistance 50Ω and gives full-scale deflection when a current of 1mA passes through it. Design a shunt that converts the galvanometer into an ammeter reading 1A at full scale. Find the resistance of the shunt and the effective resistance of the resulting ammeter.

Solution

Use the shunt formula derived above:

$$S = \frac{I_g G}{I_{max} - I_g}$$

Substitute $G = 50\Omega$, $I_g = 1\text{mA} = 0.001\text{A}$, $I_{max} = 1\text{A}$:

$$S = \frac{(0.001\text{A}) \times (50\Omega)}{1\text{A} - 0.001\text{A}} = 0.0501\Omega$$

The shunt has resistance about 0.05Ω (one twentieth of an ohm). The effective resistance of the ammeter is the parallel combination of $G = 50\Omega$ and $S = 0.05\Omega$:

$$R_{\text{ammeter}} = \frac{50\Omega \times 0.05\Omega}{50\Omega + 0.05\Omega} = 0.05\Omega$$

Making Sense of the Answer: *The completed ammeter has an effective resistance of about 0.05Ω, which is dominated by the shunt because the shunt is very much smaller than the galvanometer's 50Ω. The small effective resistance is the whole point of the design: when you insert this ammeter into a circuit to measure the current, you add only 0.05Ω of extra series resistance, which is negligible in any circuit with sensible resistor values. The ammeter is, for practical purposes, a near-invisible piece of wire. The 1000:1 ratio between the current ranges (from $1mA$ to $1A$) is achieved by a 1000:1 ratio between G and S, which is the general design rule for ammeter shunts.*

Think Like a Physicist: *The shunt formula is conservation of charge written down with Ohm's law substituted in. At full-scale deflection, I_g passes through G and the rest, $I_{max} - I_g$, passes through S; since G and S are in parallel they share the same voltage; equating the two voltages gives the formula. Never memorise the shunt formula without remembering this short derivation, because the same physics will reappear in the next example with a different geometry.*

HOT Example 24

The same galvanometer of Example 23 (50Ω resistance, 1mA full-scale deflection) is now to be converted into a voltmeter reading 10V at full scale. Find the resistance of the required multiplier, and find the effective resistance of the resulting voltmeter.

Solution

Use the multiplier formula derived earlier:

$$M = \frac{V_{max}}{I_g} - G$$

Substitute $V_{max} = 10\text{V}$, $I_g = 0.001\text{A}$, $G = 50\Omega$:

$$M = \frac{10\text{V}}{0.001\text{A}} - 50\Omega = 9950\Omega$$

The multiplier has resistance of 9950Ω.

The effective resistance of the voltmeter is the series combination of G and M:

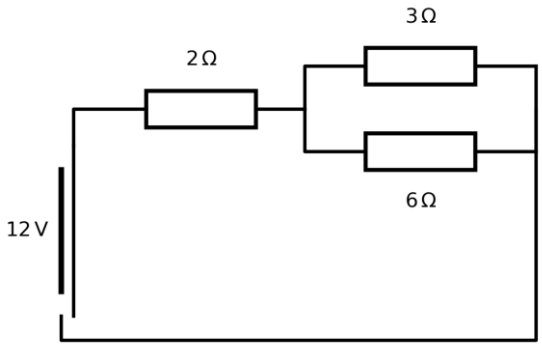
$$R_{\text{voltmeter}} = G + M = 50\Omega + 9950\Omega = 10000\Omega$$

Making Sense of the Answer: *The completed voltmeter has an effective resistance of $10k\Omega$, which is dominated by the multiplier. This is the whole point of the design: when you connect this voltmeter across a circuit element to measure its voltage, you add a parallel path of $10k\Omega$ that draws only $1mA$ of current. In most circuits this current is small enough that the voltage being measured is not appreciably disturbed. The voltmeter is, for practical purposes, a near-invisible insulator. Compare this with the ammeter of Example 23, which was a near-invisible conductor: the two instruments use the same galvanometer but achieve opposite effective resistances by adding the helper resistor in opposite places.*

Think Like a Physicist: *Every measuring instrument disturbs the circuit it measures. The design challenge in both the ammeter and the voltmeter cases is to make the disturbance small. The strategy is different for the two instruments because they measure different quantities: an ammeter is in series with the load and so must have small resistance, a voltmeter is in parallel with the load and so must have large resistance. The shunt and the multiplier are not interchangeable. Mixing them up (putting the shunt in series with the galvanometer, say) gives a useless instrument. Always draw the circuit and think about which quantity, current or voltage, the instrument is intended to measure.*

BINDER Example 25

Using the figure below, determine: (a) the equivalent resistance of the network; (b) the current through each resistor; (c) the power dissipated in each resistor.



Solution

(a) First reduce the parallel combination. With $R_1 = 3\Omega$, $R_2 = 6\Omega$:

$$R_{parallel} = \frac{R_1 R_2}{R_1 + R_2} = \frac{(3\Omega) \times (6\Omega)}{3\Omega + 6\Omega} = 2\Omega$$

Then add the 2Ω series resistor on top of this parallel combination:

$$R_{eq} = 2\Omega + 2\Omega = 4\Omega$$

(b) The current drawn from the battery is by Ohm's law applied to the equivalent resistance:

$$I = \frac{12V}{4\Omega} = 3A$$

This is the current through the 2Ω resistor (which is in series with everything else). The voltage across the parallel combination is $V = IR = (3A) \times (2\Omega) = 6V$, and this same voltage appears across both branches of the parallel pair.

So the current through the 3Ω branch is $I_3 = \frac{6V}{3\Omega} = 2A$

And the current through the 6Ω branch is $I_6 = \frac{6V}{6\Omega} = 1A$.

Check: $2A + 1A = 3A$ matches the total drawn from the battery.

(c) Power in each resistor follows from $P = I^2R$ (or equivalently $P = \frac{V^2}{R}$):

$$P_{2\Omega} = (3A)^2 \times 2\Omega = 18W$$
$$P_{3\Omega} = (2A)^2 \times 3\Omega = 12W$$
$$P_{6\Omega} = (1A)^2 \times 6\Omega = 6W$$

Making Sense of the Answer: *Total power dissipated:* $18W + 12W + 6W = 36W$. *Check against* $P = IV = (3A) \times (12V) = 36W$ *of total power delivered by the battery. The two numbers agree, as they must by conservation of energy. The* 2Ω *resistor (which carries the full current) gets the largest share of the power; the* 6Ω *branch (which carries the smallest current) gets the smallest share. Both observations are general: in a series chain, the largest resistance dissipates the most; in a parallel cluster, the smallest resistance dissipates the most.*

Think Like a Physicist: *Mixed networks reduce by working from the innermost knot outwards. Spot the parallel block first, replace it by its equivalent, then the rest of the network is a simple series chain. Once the current through the equivalent resistor is known, work back into the parallel block to find how the current divides between its branches. The two-step strategy (outwards to find the total, inwards to find the parts) works for every mixed network this section can handle.*

REAL Example 26

A 6V torch is powered by four identical zinc-carbon cells connected in series, each of EMF 1.5V and internal resistance 0.2Ω. The torch bulb has resistance 4Ω when hot. Find the total EMF of the battery stack, the total internal resistance, the current that flows through the bulb, and the power dissipated in the bulb.

Solution

Four cells in series each contribute their EMF and internal resistance to the totals. Using the series-cells formulas:

$$E_{total} = nE = 4 \times 1.5V = 6V$$
$$r_{total} = nr = 4 \times 0.2\Omega = 0.8\Omega$$

The current through the bulb is determined by the total EMF, which is given by:

$$E_{total} = I(R + r_{total})$$
$$I = \frac{E_{total}}{R + r_{total}} = \frac{6V}{4\Omega + 0.8\Omega} = 1.25A$$

Power dissipated in the bulb is $P = I^2R$:

$$P = (1.25A)^2 \times 4\Omega = 6.25W$$

Making Sense of the Answer: *The bulb dissipates* $6.25W$*, which is a reasonable figure for a torch bulb. Some of the available electrical power is also lost inside the cells:* $P_{internal} = I^2 r_{total} = 1.25^2 \times 0.8 = 1.25W$*, which warms the cells while the torch is on. The total power drawn from the battery stack is* $6.25W + 1.25W = 7.5W$*, which matches* $I \times E_{total} = 1.25A \times 6V = 7.5W$ *as it must.*

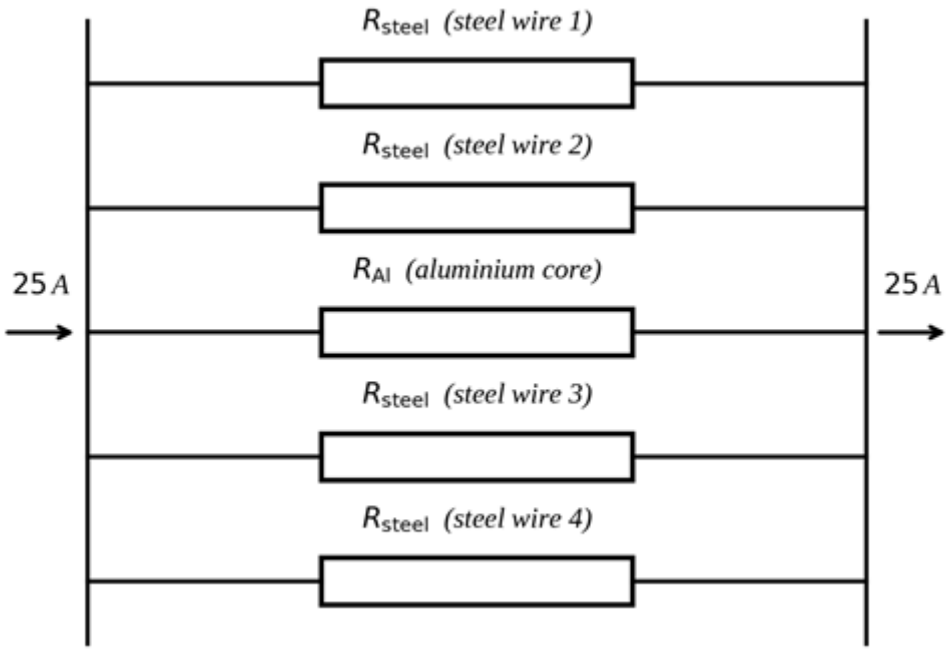
Think Like a Physicist: *Two facts about this torch are worth noticing. First,* 17% *of the chemical energy of the cells is lost as heat inside the cells themselves (the ratio* $\frac{r_{total}}{R+r_{total}}$ *of internal to total resistance equals the fraction of power wasted internally). This is an inevitable loss; nothing about circuit design can recover it. Second, as the cells age, their internal resistance* r *grows. The total internal resistance grows with it, and a larger fraction of the total power is wasted inside the cells, leaving less for the bulb. This is why old torch cells produce a dim light: not because the EMF has dropped much, but because too much of the available power is now being wasted internally.*

HOT Example 27

A composite TANESCO transmission cable consists of an aluminium core surrounded by four steel reinforcing wires, each of the same length and the same cross-sectional area as the aluminium core. The total current carried by the cable is 25A. The resistivities are $\rho_{Al} = 2.82 \times 10^{-8}\Omega m$ and $\rho_{steel} = 1.6 \times 10^{-7}\Omega m$. Find the current carried by the aluminium core and the current carried by each of the four steel wires.

Solution

The five wires (one aluminium core, four steel reinforcement wires) run the whole length of the cable and meet at the cable's two ends. So all five share the same pair of endpoints, which means they are connected in parallel. The following circuit diagram sets the parallel network out explicitly.



Since all five wires have the same length and the same cross-section, their geometric factor l/A is the same for each wire, and hence the resistance R = ρl/A is proportional to the resistivity ρ alone. A ratio of two resistances therefore reduces to a ratio of two resistivities.

Computing that ratio:

$$\frac{R_{steel}}{R_{Al}} = \frac{\rho_{steel}}{\rho_{Al}} = \frac{1.6 \times 10^{-7}\Omega m}{2.82 \times 10^{-8}\Omega m} = 5.67 \text{ or } R_{steel} = 5.67R_{Al}$$

Since the five wires are in parallel, the voltage across each of them is the same value V. The current in each wire is given by Ohm's law applied to that wire alone:

$$I_{Al} = \frac{V}{R_{Al}}, \quad I_{steel} = \frac{V}{R_{steel}}$$

Taking the ratio of these two equations cancels V and gives the current ratio in terms of the resistance ratio:

$$\frac{I_{steel}}{I_{Al}} = \frac{R_{Al}}{R_{steel}} = \frac{R_{Al}}{5.67R_{Al}} = \frac{1}{5.67} = 0.176 \text{ or } I_{steel} = 0.176I_{Al}$$

The total current is the sum across all five wires:

$$25A = I_{Al} + 4 \times I_{steel} = I_{Al} + 4 \times (0.176I_{Al}) = 1.706I_{Al}$$

$$I_{Al} = 14.7A$$

And the current in each steel wire follows:

$$I_{steel} = 0.176 \times 14.7A = 2.6A$$

Checking the totals: aluminium 14.7A plus four steel wires at 2.6A each gives $14.7A + 4 \times 2.6A = 25.1A$, which agrees with the given total of 25A to within the rounding.

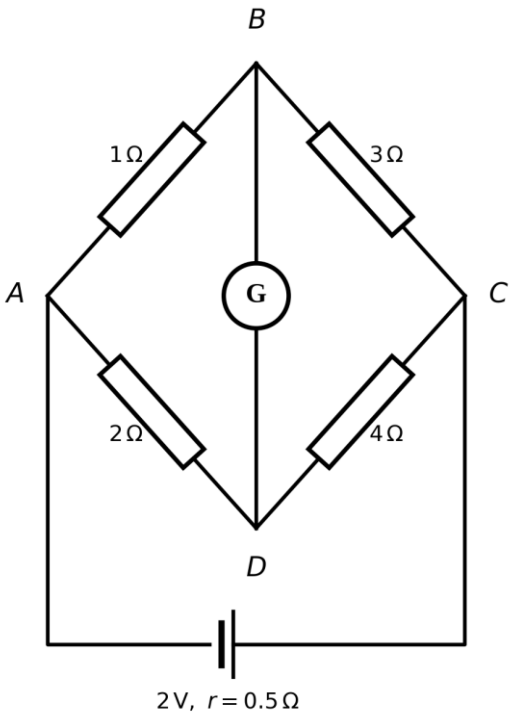
Making Sense of the Answer: *The aluminium core carries* 14.7*A out of the total* 25*A, which is about* 59% *of the cable's current. The four steel wires together carry the remaining* 41%*, split equally as* 2.6*A each. The aluminium is a single wire out of five and carries more current than all four steel wires combined; this is because aluminium is* 5.67 *times more conductive than steel, and the steel wires divide the load equally among themselves, four ways.*

The Physics Behind the Design: *The steel wires are not in the cable to carry current efficiently. They are in the cable to keep it from stretching under its own weight as it hangs between distant pylons. Aluminium is a soft metal that creeps and sags; steel is strong but a poor conductor. The composite cable combines the two properties: the aluminium does most of the conducting and the steel does most of the holding. Engineering trade-offs of exactly this kind are why real transmission cables are rarely made of a single material.*

Think Like a Physicist: *When a parallel-resistor problem has all the geometric factors (length, area, shape) the same across every branch, do not compute the absolute resistances. Work with the ratio of the remaining quantity (here the resistivity) and let it carry the whole calculation. The geometric factors cancel out of every ratio, and the arithmetic stays close to numbers of order unity rather than scientific notation. This habit, working with ratios rather than absolute values whenever the problem allows, is one of the most powerful tools a physicist has, and not just in circuit problems.*

HOT Example 28

Four resistors of resistances 1Ω, 3Ω, 4Ω, and 2Ω are connected along the four sides of a quadrilateral ABCD in that order (so AB = 1Ω, BC = 3Ω, CD = 4Ω, DA = 2Ω). An accumulator of EMF 2V and internal resistance 0.5Ω is connected across A and C. A galvanometer of resistance 5Ω is connected between B and D. Find the current through the galvanometer and the effective resistance of the network between A and C.



Solution

Try first the methods of this section. The four resistors form two paths from A to C: a top path through B (with AB = 1Ω then BC = 3Ω in series, total 4Ω) and a bottom path through D (with AD = 2Ω then DC = 4Ω in series, total 6Ω). If the network had no galvanometer between B and D, the two paths would be a simple parallel pair, and the equivalent resistance between A and C would be (4Ω)(6Ω)/(4Ω + 6Ω) = 2.4Ω.

But the galvanometer disturbs this simple picture. The resistor ratio AB:BC is 1:3, and the resistor ratio AD:DC is 2:4 = 1:2. The two ratios are **not** equal. So the potential at B (which is shared between AB and BC in ratio 3:1 of the A-to-C voltage drop) is **not** equal to the potential at D (which is shared in ratio 2:1 of the same drop). Because the potentials at B and D are

different, current flows through the galvanometer connecting them, and this current alters how the original two paths share the total.

This is the unbalanced Wheatstone bridge, in its simplest form. It cannot be reduced by series and parallel alone (the galvanometer branch creates a fifth connection that breaks the simple parallel topology). It does not have the symmetry of the cube either (the four resistor values are all different from one another). The methods of this section are therefore not enough.

The unbalanced bridge yields cleanly to Kirchhoff's laws, which the next section introduces. We shall return to the exact numerical solution after Kirchhoff is in our hands. For now, note that the diagnosis (unbalanced bridge, not reducible by the methods of this section) is the most important conceptual step. Identifying the right tool for the job is half the physics.

Making Sense of the Anticipation: *If the four resistor values had been in balanced ratio AB:BC = AD:DC (for example, 1Ω, 2Ω, 2Ω, 4Ω rather than 1Ω, 3Ω, 4Ω, 2Ω), the potentials at B and D would have been equal, no current would flow through the galvanometer, and the network would have reduced to a simple parallel pair. The fact that the actual values are unbalanced is what makes the problem non-trivial. The balanced case (galvanometer reads zero) is the working principle of the Wheatstone bridge as a measurement device, and we shall meet it again in a later section.*

Think Like a Physicist: *Recognising which method applies to which network is the most important skill of this section. Use series and parallel reductions whenever they work; they always work fastest. Use symmetry methods when geometric symmetry forces equipotentials. Use Kirchhoff's laws (the next section) as a fallback for everything else. The fallback is the most general tool, but it is also the slowest, so never reach for it first.*

Seven examples and one escape. Series, parallel, symmetry, and meter modification carried us through the first six. The last one stayed beyond our reach. For that we need a sharper tool.

That tool is Kirchhoff's laws, two short statements about conservation in circuits that solve every network in principle by reducing it to a set of simultaneous linear equations. The next section sets them out and then applies them in earnest.

KIRCHHOFF'S LAWS

Every method in this chapter so far has had a limited reach. Ohm's law on its own handles a single resistor. Series and parallel reduction handles any network that can be untangled into nested series and parallel pieces. Symmetry methods handle the special networks whose geometry forces certain nodes to the same potential. Each method is fast within its own territory, and each fails outside it. The unbalanced bridge that closed the last section defeated all three.

This section introduces the method that has no such limit. Kirchhoff's two laws apply to every network without exception, however many sources it contains, however the resistors are wired, and whether or not it has any symmetry. The price for this generality is labour: Kirchhoff's laws turn a circuit into a set of simultaneous equations that must then be solved. They are the method of last resort, to be reached for only when the faster methods fail, but they never fail themselves.

Both laws are conservation statements that you already know from earlier physics, applied now one piece of a circuit at a time. The first law is conservation of charge applied at a junction. The second law is conservation of energy applied around a loop. Before stating them, we need three pieces of vocabulary.

Junctions, branches, and loops

A junction (also called a **node**) *is a point in a circuit where three or more wires meet.* At a junction the current has a choice of paths, so the current can **split** or **merge** there. A point where only two wires meet is **not** a junction; it is just a point along a single wire, and the same current flows through it on both sides.

A branch *is a single path between two junctions, along which the current is the same all the way.* A branch may contain several components in series (resistors, cells, or both), but because it has no junction in its interior, the current cannot change anywhere along it. Each branch carries one current, and that current is one of the unknowns we shall solve for.

A loop *is any closed conducting path in the network:* a path that starts at one point, travels through some branches, and returns to its starting point without retracing any branch. A network with several junctions usually has several possible loops, and part of the skill of applying Kirchhoff's laws is choosing a convenient set of them.

Kirchhoff's current law

Kirchhoff's current law, abbreviated **KCL**, states that: *The algebraic sum of the currents at any junction is zero.* Said more plainly: *the total current flowing into a junction equals the total current flowing out of it.* Whatever charge arrives at the junction each second must leave it again the same second.

The reason is conservation of charge. A junction is just a point, and a point has no room to store charge. Charge cannot pile up there and cannot be created or destroyed there, so every coulomb that flows in must immediately flow out along one path or another. If this were not so, charge would accumulate at the junction without limit, which never happens in a steady circuit.

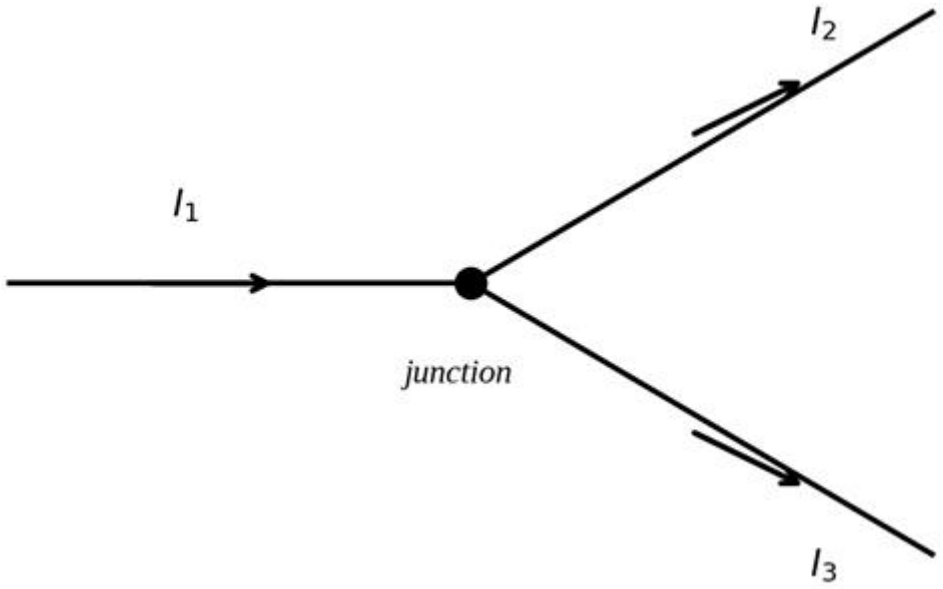


Figure: ***A junction where three wires meet.*** *Current I_1 flows in; currents I_2 and I_3 flow out. Kirchhoff's current law requires $I_1 = I_2 + I_3$, the current in equals the current out.*

For the junction in the figure, with I_1 flowing in and I_2 and I_3 flowing out, the law reads:

$$\mathbf{I_1 = I_2 + I_3}$$

To apply the law systematically, adopt a sign convention: *count currents flowing into the junction* as **positive** and *currents flowing out* as **negative**. The algebraic sum is then zero. With this convention the same junction reads:

$$\mathbf{I_1 - I_2 - I_3 = 0},$$

which rearranges to the same statement. The two forms are identical; use whichever is clearer for the problem at hand.

Kirchhoff's voltage law

Kirchhoff's voltage law, abbreviated **KVL**, states that *around any closed loop in a network, the algebraic sum of the EMFs equals the algebraic sum of the potential drops IR.* In symbols, taken around one complete loop:

$$\mathbf{\Sigma E = \Sigma IR}$$

The reason is conservation of energy. Follow one coulomb of charge once around the loop and back to where it started. Along the way the EMF sources do work on the coulomb (raising its energy) and the resistors take energy from it (lowering its energy as heat). Since the coulomb returns to exactly where it began, its energy must be back to its starting value, so the total energy gained from the sources must equal the total energy lost in the resistors. Energy per coulomb is voltage, so the sum of the EMFs equals the sum of the IR drops.

This is the same coulomb-tracing argument used earlier in the chapter to derive the terminal-voltage relation for a single cell. Kirchhoff's voltage law is that argument promoted to a general principle and applied to any loop, not just a single-cell circuit.

Applying the law requires a sign convention, because both the EMFs and the IR drops can act in either direction around the loop. Begin by choosing a direction to traverse the loop, either clockwise or anticlockwise; the choice is free, but once made it must be held for the whole loop. Then:

*An **EMF** counts as positive when the chosen loop direction passes through the source from its negative terminal to its positive terminal (the direction in which the source raises the potential). It counts as negative when the loop direction passes through it the other way.*

*An **IR drop** counts as positive when the chosen loop direction matches the assumed direction of the current through that resistor. It counts as negative when the loop direction opposes the assumed current. The product IR uses the branch current I with its sign, so a branch current that later turns out negative takes care of itself automatically.*

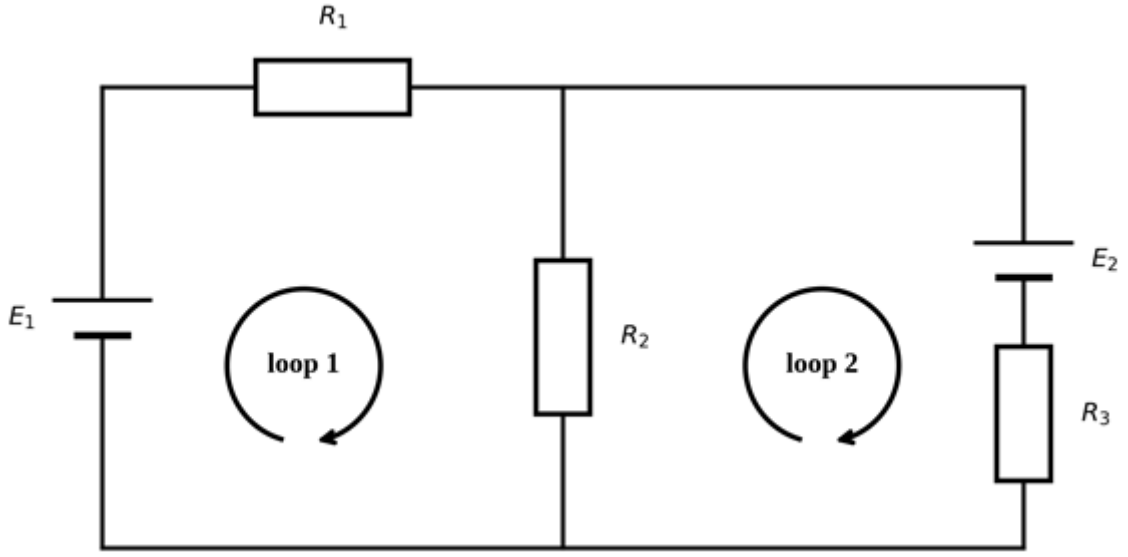


Figure*: A two-loop network.* *Each loop is traversed in a chosen direction (here both clockwise, shown by the circular arrows). Kirchhoff's voltage law is written once for each loop, and Kirchhoff's current law once for each independent junction. Together they give enough equations to solve for all the branch currents.*

The procedure, step by step

Solving any network by Kirchhoff's laws follows four steps. Each is explained below with what to do in practice.

Step 1: Label every branch current with an assumed direction

Mark a current symbol and an arrow on every branch. The direction of each arrow is a guess; you do not need to guess correctly. If a current turns out to flow the other way, its value will simply come out negative at the end, and that negative sign is the circuit telling you the true direction is opposite to your guess. Use Kirchhoff's current law at the junctions immediately to reduce the number of distinct unknowns: if three branches meet at a junction, the third current is fixed once the first two are named.

Step 2: Write the current law at each independent junction

A network with **n** junctions gives $\mathbf{n-1}$ independent current-law equations. (The last junction gives no new information, because charge conservation across the whole network makes it follow automatically from the others.) Write the current law at all but one of the junctions.

Step 3: Write the voltage law around enough independent loops

Choose a set of loops so that every branch is included in at least one loop, and write Kirchhoff's voltage law for each, using the sign convention above. You need exactly as many loop equations as there are remaining unknown currents after the junction equations have been used. A loop is independent if it includes at least one branch that no previous loop has used.

Step 4: Solve the simultaneous equations

The junction and loop equations together form a system of linear simultaneous equations in the unknown currents. Solve them. A negative answer for any current means the actual current in that branch flows opposite to the direction you assumed in Step 1; the magnitude is still correct. Always finish by checking the solution in one equation you did not use to find it, or by confirming that the power delivered by the sources equals the power dissipated in the resistors.

Parallel cells with different EMFs

We can now settle a question left open earlier in the chapter: *what happens when two cells of different EMF are connected in parallel across a load?* The equal-EMF case was simple, but the unequal case needs Kirchhoff's laws, and it makes an instructive first application of the procedure.

Two cells are connected in parallel across a load resistor R. The first cell has EMF E_1 and internal resistance r_1; the second has EMF E_2 and internal resistance r_2. The figure shows the circuit, with the cell currents I_1 and I_2 both assumed to flow out of their cells toward the load, and the load current I flowing down through R.

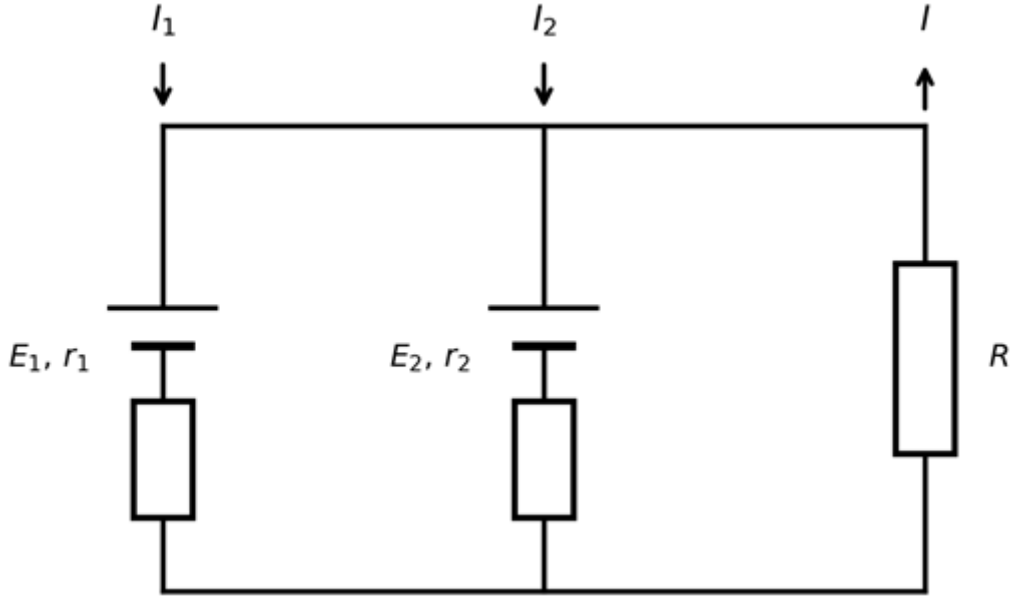


Figure: ***Two cells of different EMF in parallel, driving a common load R.*** *Cell 1 supplies current I_1, cell 2 supplies current I_2, and the load carries $I = I_1 + I_2$.*

Apply the current law (KCL) at the top junction, where both cell currents meet and the load current leaves:

$$I = I_1 + I_2$$

Apply the voltage law (KVL) around the loop containing cell 1 and the load. The EMF E_1 drives the current; it is spent on the internal-resistance drop $I_1 r_1$ inside cell 1 and the load drop IR across R:

$$E_1 = I_1 r_1 + IR$$

Apply the voltage law around the loop containing cell 2 and the load:

$$E_2 = I_2 r_2 + IR$$

These are three equations in the three unknowns I_1, I_2, and I. Solve them.

From the two loop equations, express each cell current in terms of I:

$$I_1 = \frac{E_1 - IR}{r_1}, \quad I_2 = \frac{E_2 - IR}{r_2}$$

Substitute both into the junction equation $I = I_1 + I_2$ and collect the terms in I:

$$I = \frac{E_1 - IR}{r_1} + \frac{E_2 - IR}{r_2} = \frac{r_2(E_1 - IR) + r_1(E_2 - IR)}{r_1 r_2} = \frac{E_1 r_2 + E_2 r_1 - IR(r_2 + r_1)}{r_1 r_2}$$

$$Ir_1 r_2 = E_1 r_2 + E_2 r_1 - IR(r_2 + r_1);\ I\big(r_1 r_2 + R(r_2 + r_1)\big) = E_1 r_2 + E_2 r_1$$

Hence:

$$\mathbf{I = \frac{E_1 r_2 + E_2 r_1}{R(r_1 + r_2) + r_1 r_2}}$$

This is the load current for any two cells in parallel. The formula is worth reading for its structure rather than memorising. Compare it with the single-source result, load current equals EMF over (load plus internal resistance), and it suggests that the parallel pair behaves like one equivalent cell. Indeed it does.

Rewriting the result in the single-source form $I = \frac{E_{eq}}{R + r_{eq}}$ identifies the equivalent EMF and equivalent internal resistance of the parallel pair:

$$\mathbf{E_{eq} = \frac{E_1 r_2 + E_2 r_1}{r_1 + r_2}}, \quad \mathbf{r_{eq} = \frac{r_1 r_2}{r_1 + r_2}}$$

The equivalent internal resistance is just r_1 and r_2 in parallel, as expected for two resistances side by side. The equivalent EMF is a weighted average of the two EMFs, each weighted by the other cell's internal resistance.

Check the formula against the equal-cell case. Put $E_1 = E_2 = E$ and $r_1 = r_2 = r$. Then $E_{eq} = \frac{Er + Er}{2r} = E$ and $r_{eq} = \frac{r^2}{2r} = \frac{r}{2}$, which recovers exactly the equal-cell result found earlier in the chapter. The general formula contains the simple one as a special case, which is a good sign that it is right.

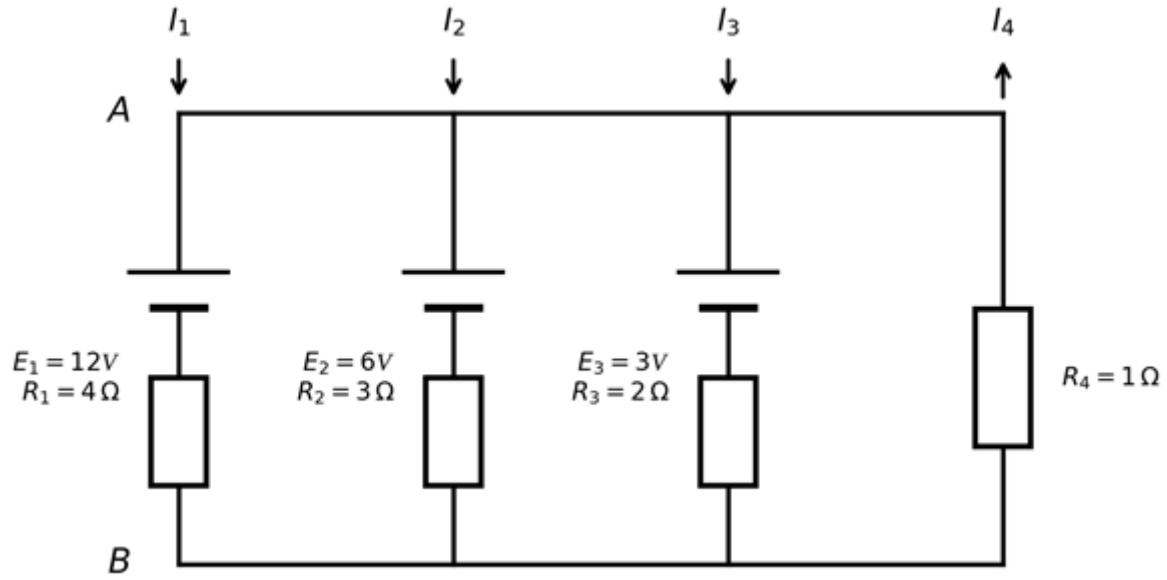
A practical warning: *Connecting cells of different EMF in parallel is generally a bad idea. Even with no external load connected, the two cells form a closed loop by themselves, and the higher-EMF cell drives a current backwards through the lower-EMF cell, charging it and wasting energy as heat in both internal resistances. Worked Example 29 shows this directly:*

one of the two cell currents comes out negative, meaning that cell is being charged by the other rather than helping to drive the load. Parallel cells should have matched EMFs.

Five examples follow. Each is solved by the same four-step procedure, and several contain a current that comes out negative, so watch how the sign is read at the end.

HOT Example 29

Three cells are connected in parallel with a load resistor across a pair of common rails, as shown. The cells are $E_1 = 12V$ with $R_1 = 4\Omega$, $E_2 = 6V$ with $R_2 = 3\Omega$, and $E_3 = 3V$ with $R_3 = 2\Omega$. The load is $R_4 = 1\Omega$. Find the current in each of the four branches.



Solution

There are only two junctions, the top rail A and the bottom rail B. Let V be the potential difference between them, $V = V_A - V_B$. This single unknown is the key: once V is known, every branch current follows from Ohm's law applied to that branch.

For each cell branch, the terminal voltage equals the EMF minus the internal drop:

$$V = E_i - I_i R_i \text{ or } I_i = \frac{(E_i - V)}{R_i}$$

For the load branch:

$$I_4 = \frac{V}{R_4}$$

The KCL at A says the three cell currents flowing in equal the load current flowing out:

$$\frac{E_1 - V}{R_1} + \frac{E_2 - V}{R_2} + \frac{E_3 - V}{R_3} = \frac{V}{R_4}$$

$$\frac{12 - V}{4} + \frac{6 - V}{3} + \frac{3 - V}{2} = \frac{V}{1}$$

Solving for V:

$$V = 3.12V$$

Now back-substitute to get each branch current:

$$I_1 = \frac{12V - 3.12V}{4\Omega} = 2.22A$$

$$I_2 = \frac{6V - 3.12V}{3\Omega} = 0.96A$$

$$I_3 = \frac{3V - 3.12V}{2\Omega} = -0.06A$$

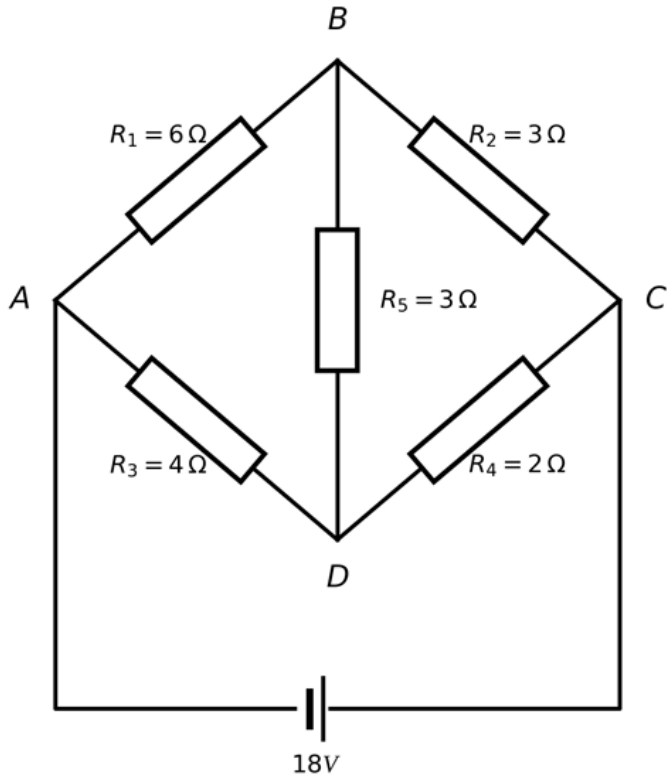
$$I_4 = \frac{3.12V}{1\Omega} = 3.12A$$

Making Sense of the Answer: ***Check the current law:*** $2.22A + 0.96A + (-0.06A) = 3.12A$*, which equals* I_4*, as it must. The interesting result is* I_3*, which came out negative. The assumed direction for* I_3 *was upward, out of the 3V cell. The negative sign means the actual current is downward: the 3V cell is too weak to push against the common rail voltage of 3.12V, so instead of delivering current it receives a small current of 0.06A and is being charged by the other two cells. This is exactly the behaviour the parallel-cell warning predicted.*

Think Like a Physicist: *This network had only two junctions, so a single unknown (the rail voltage* V*) captured the whole problem. Whenever a network reduces to two nodes joined by several parallel branches, the node-voltage shortcut used here is faster than writing a separate current symbol for every branch and solving the full system. Recognising when a circuit has this two-node form can save a great deal of algebra.*

HOT Example 30

In the bridge network shown, four resistors form a diamond ABCD: $R_1 = 6\Omega$ on AB, $R_2 = 3\Omega$ on BC, $R_3 = 4\Omega$ on AD, and $R_4 = 2\Omega$ on DC. A fifth resistor $R_5 = 3\Omega$ bridges B to D. A 18V source of negligible internal resistance is connected across A and C. Find the current in all five branches.



Solution

Before grinding through Kirchhoff's laws, look at the arm resistances. Going down the upper path A → B → C, the ratio of the two arms is $\frac{R_1}{R_2} = \frac{6\Omega}{3\Omega} = 2$. Going down the lower path A → D → C, the ratio is $\frac{R_3}{R_4} = \frac{4\Omega}{2\Omega} = 2$. The two ratios are equal. This is the balance condition of the bridge, and it has a striking consequence.

When the arm ratios are equal, the potentials at B and D are equal. To see why, note that the upper path divides the full A-to-C voltage between R_1 and R_2 in the ratio 6 to 3, while the lower path divides the same voltage between R_3 and R_4 in the ratio 4 to 2. Both ratios are 2 to 1, so the point B sits at the same fraction of the way down from A to C as the point D does. They are at the same potential.

Since $V_B = V_D$, the voltage across the bridge resistor R_5 is zero, and so by Ohm's law the current through it is zero:

$$I_5 = \frac{V_B - V_D}{R_5} = 0$$

With no current in the bridge, the network becomes two simple series paths in parallel. The upper path has resistance $R_1 + R_2 = 6\Omega + 3\Omega = 9\Omega$, carrying current $I_{upper} = \frac{18V}{9\Omega} = 2A$.

The lower path has resistance $R_3 + R_4 = 4\Omega + 2\Omega = 6\Omega$, carrying current $I_{lower} = \frac{18V}{6\Omega} = 3A$.

So the five branch currents are: 2A through R_1, 2A through R_2 (the whole upper path carries the same 2A), 3A through R_3, 3A through R_4 (the whole lower path carries 3A), and zero through the bridge resistor R5.

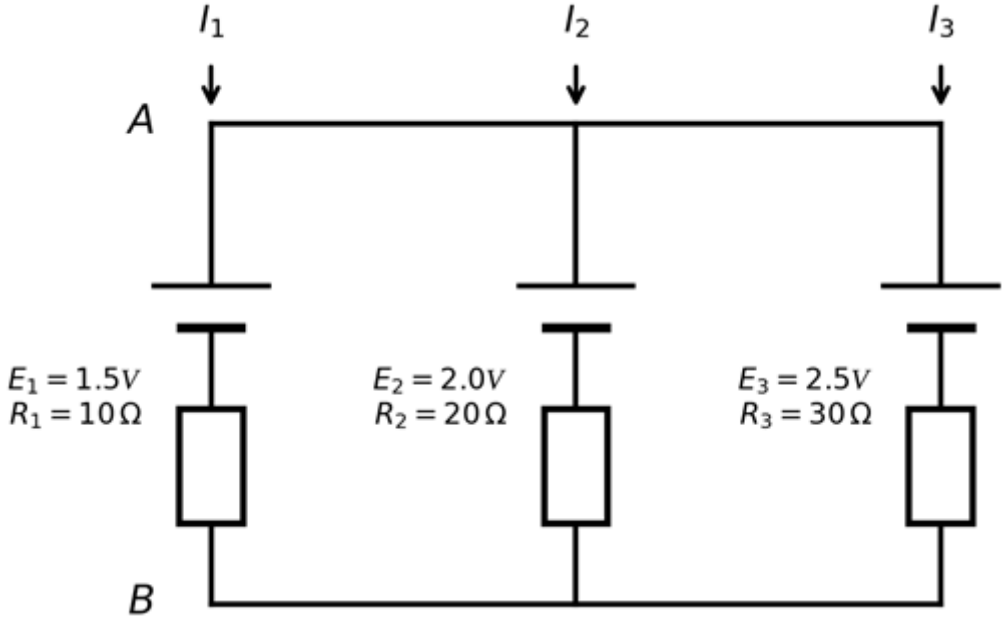
Making Sense of the Answer: *The total current drawn from the source is $2A + 3A = 5A$, which leaves the source, splits at A into the two paths, and recombines at C. The zero bridge current is the heart of the Wheatstone bridge as a measuring instrument: when the arms are balanced, a galvanometer placed across BD reads exactly zero, and that null reading can be detected with great precision. The next section builds an entire measurement technique on this fact.*

Think Like a Physicist: *Always test a bridge for balance before reaching for the full Kirchhoff machinery. The balance check is one division: compare $\frac{R_1}{R_2}$ with $\frac{R_3}{R_4}$. If they are equal, the bridge branch carries no current and the network collapses*

to two series paths in parallel, which you can solve in your head. Only when the ratios differ (the unbalanced case of Example 33) do you need the full simultaneous equations.

HOT Example 31

Three cells of negligible internal resistance are connected in parallel between two points A and B, each cell in series with its own resistor: $E_1 = 1.5V$ with $R_1 = 10\Omega$, $E_2 = 2.0V$ with $R_2 = 20\Omega$, and $E_3 = 2.5V$ with $R_3 = 30\Omega$. All three positive terminals face A. Find the current in each resistor.



Solution

As in Example 29, there are only two junctions, so let $V = V_A - V_B$ be the single unknown. Each branch current, assumed to flow from B up to A, is $I_i = (E_i - V)/R_i$.

Since A and B are the only nodes and these three branches are the only paths between them, the current law requires the three currents into A to sum to zero:

$$\frac{1.5 - V}{10} + \frac{2.0 - V}{20} + \frac{2.5 - V}{30} = 0$$

Solve for V:

$$V = 1.818V$$

Substitute back to find each current:

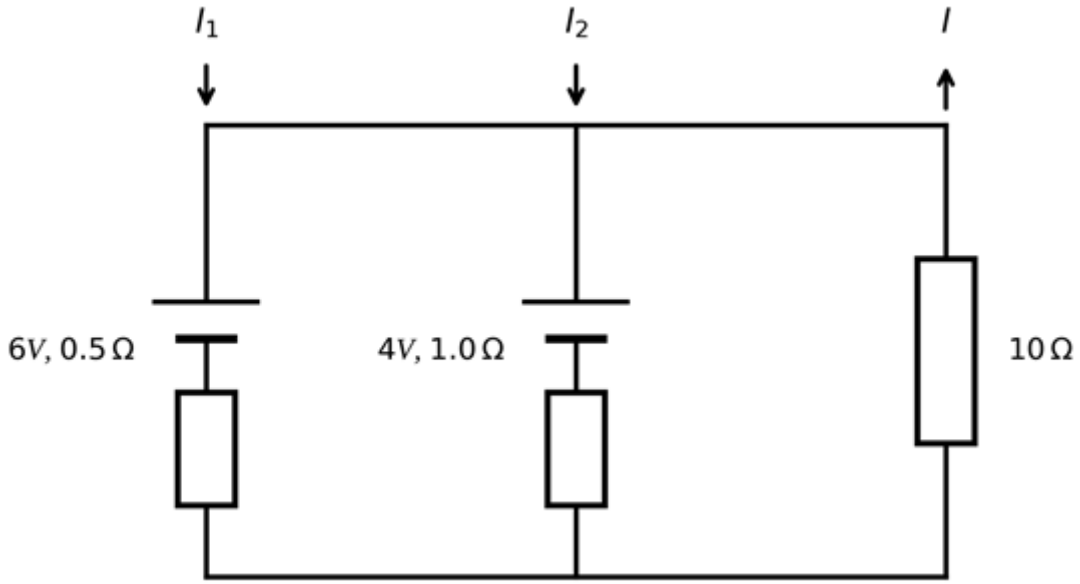
$$I_1 = \frac{1.5V - 1.818V}{10\Omega} = -0.032A = -32mA$$
$$I_2 = \frac{2.0V - 1.818V}{20\Omega} = 0.0091A = 9.1mA$$
$$I_3 = \frac{2.5V - 1.818V}{30\Omega} = 0.0227A = 22.7mA$$

Making Sense of the Answer: ***Check:*** $-32mA + 9.1mA + 22.7mA \approx 0mA$*, confirming the current law. The current* I_1 *came out negative. The 1.5V cell has the lowest EMF of the three, lower than the common rail voltage of 1.818V that the other two establish, so it cannot push current out into the rail. Instead it absorbs current: the actual flow in branch 1 is downward at 32mA, and the 1.5V cell is being charged by the 2.0V and 2.5V cells. The higher-EMF cells, being above the rail voltage, deliver current as expected.*

Think Like a Physicist: *The common rail voltage 1.818V is a kind of compromise between the three EMFs, pulled toward the cells with the smaller series resistance (which have more influence). Any cell whose EMF lies above this compromise voltage delivers current; any cell whose EMF lies below it absorbs current. This is the general fate of mismatched cells wired in parallel, and the reason the practice is discouraged.*

REAL Example 32

Two cells are connected in parallel across a 10Ω load. The first cell has $E_1 = 6V$, $r_1 = 0.5\Omega$; the second has $E_2 = 4V$, $r_2 = 1.0\Omega$. Find (a) the load current; (b) the current through each cell; (c) the equivalent EMF and internal resistance of the parallel combination.



Solution

(a) Start with the equivalent source, using the formulas derived in this section. The equivalent internal resistance is r_1 and r_2 in parallel:

$$r_{eq} = \frac{r_1 r_2}{r_1 + r_2} = \frac{0.5\Omega \times 1.0\Omega}{0.5\Omega + 1.0\Omega} = 0.333\Omega$$

The equivalent EMF is the weighted average:

$$E_{eq} = \frac{E_1 r_2 + E_2 r_1}{r_1 + r_2} = \frac{6V \times 1.0\Omega + 4V \times 0.5\Omega}{0.5\Omega + 1.0\Omega} = 5.33V$$

(b) The load current follows from the equivalent source driving the 10Ω load:

$$I = \frac{E_{eq}}{R + r_{eq}} = \frac{5.33V}{10\Omega + 0.333\Omega} = 0.516A$$

(c) The voltage across the load is $V = IR = 0.516A \times 10\Omega = 5.16V$. This same voltage appears across each cell's terminals (they are all in parallel). The current from each cell is then $I_i = (E_i - V)/r_i$:

$$I_1 = \frac{6V - 5.16V}{0.5\Omega} = \frac{0.84}{0.5} = 1.68A$$

$$I_2 = \frac{4V - 5.16V}{1.0\Omega} = \frac{-1.16}{1.0} = -1.16A$$

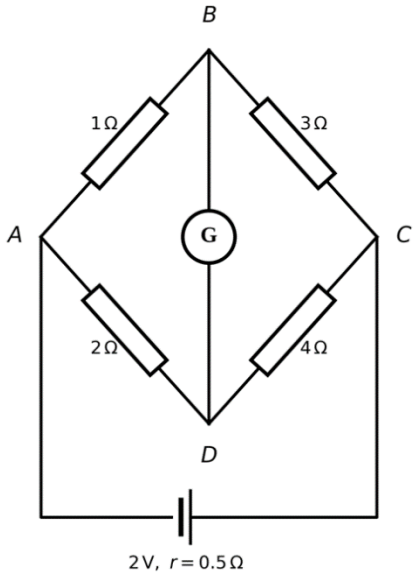
Making Sense of the Answer: ***Check the current law:*** $1.68 + (-1.16) = 0.52A$, *which matches the load current* $0.516A$ *to within rounding. The second cell's current is negative. Its EMF of 4V is below the common terminal voltage of 5.16V that the stronger 6V cell establishes, so the 4V cell cannot deliver current. Instead it absorbs 1.16A: it is being charged by the 6V cell. The 6V cell is supplying 1.68A, of which 0.52A goes to the load and 1.16A goes uselessly into charging its partner. This is the concrete cost of mismatched parallel cells: more than two thirds of the strong cell's output is wasted on the weak one.*

Think Like a Physicist: *The equivalent-source method turned a three-unknown problem into two short formulas plus one Ohm's-law division. Once the parallel pair is replaced by its single equivalent EMF and internal resistance, the rest of the circuit sees nothing but an ordinary cell. Replacing a complicated sub-network by its equivalent source is one of the most useful simplifications in circuit analysis, and it is worth doing whenever a group of sources and resistors feeds the rest of the circuit through a single pair of terminals.*

We can now finish the problem left unsolved at the end of the previous section.

HOT Example 33

Four resistors lie along the sides of a quadrilateral ABCD: $AB = 1\Omega$, $BC = 3\Omega$, $CD = 4\Omega$, $DA = 2\Omega$. A galvanometer of resistance 5Ω connects B to D. An accumulator of EMF 2V and internal resistance 0.5Ω is connected across A and C. Find the galvanometer current and the resistance of the network between A and C.



Solution

First confirm that the bridge is unbalanced, so that the shortcut of Example 30 does not apply. The upper ratio is $\frac{AB}{BC} = \frac{1}{3}$; the lower ratio is $\frac{AD}{DC} = \frac{2}{4} = \frac{1}{2}$. These are not equal, so $V_B \neq V_D$ and current does flow through the galvanometer. We use the node-voltage form of Kirchhoff's laws.

Take C as the zero of potential ($V_C = 0$). The unknowns are the potentials V_A, V_B, V_D.

Write the current law at each of the three nodes, expressing every branch current as a potential difference over a resistance.

At node A, the accumulator delivers current $(2 - V_A)/0.5$ into the node, and this leaves through AB and AD:

$$\frac{2 - V_A}{0.5} = \frac{V_A - V_B}{1} + \frac{V_A - V_D}{2}$$

$$7V_A - 2V_B - V_D = 8 \ldots \text{(i)}$$

At node B, current arrives from A and leaves through BC and the galvanometer BD:

$$\frac{V_A - V_B}{1} = \frac{V_B}{3} + \frac{V_B - V_D}{5}$$

$$15V_A - 23V_B + 3V_D = 0 \ldots \text{(ii)}$$

At node D, current arrives from A and from the galvanometer, and leaves through DC:

$$\frac{V_A - V_D}{2} + \frac{V_B - V_D}{5} = \frac{V_D}{4}$$

$$10V_A + 4V_B - 19V_D = 0 \ldots \text{(iii)}$$

These are three linear equations in the three unknown potentials. Solving with a calculator gives:

$$V_A = 1.654\text{V}, \quad V_B = 1.226\text{V}, \quad V_D = 1.129\text{V}$$

The galvanometer current flows from B to D (since $V_B > V_D$):

$$I_g = \frac{V_B - V_D}{5} = \frac{1.226\text{V} - 1.129\text{V}}{5\Omega} = 0.0195\text{A} = 19.5\text{mA}$$

For the resistance between A and C, first find the total current the accumulator drives:

$$I = \frac{2 - V_A}{0.5} = \frac{2\text{V} - 1.654\text{V}}{0.5\Omega} = 0.692\text{A}$$

The network resistance between A and C (not counting the accumulator's own internal resistance) is the terminal voltage divided by this current:

$$R_{AC} = \frac{V_A}{I} = \frac{1.654\text{V}}{0.692\text{A}} = 2.39\Omega$$

Making Sense of the Answer: *The galvanometer carries about* $19.5mA$*, a small but definitely non-zero current, exactly as the unbalanced arm ratios predicted. Had the bridge been balanced, this current would have been zero and the network resistance would have been the simple parallel combination of the two arm paths. The unbalanced current is what makes this problem need the full Kirchhoff treatment. As a sanity check,* $I = 2/(R_{AC} + 0.5) = 2/2.89 = 0.692A$*, which matches the total current found above.*

Think Like a Physicist: *The node-voltage method used here is often the quickest form of Kirchhoff's laws for a network with few nodes. Instead of naming a separate current for every branch and juggling both laws, name the node potentials, write the*

current law once per node with each current expressed as a voltage difference over a resistance, and the voltage law takes care of itself. Three nodes gave three equations and the whole bridge fell out. The branch-current method and the node-voltage method always agree; choose whichever gives fewer unknowns for the network in front of you.

Five worked examples and five networks now lie behind us.

One particular network in this section, the balanced bridge of Example 30, hid a special gift: when its arms are in equal ratio, the bridge branch carries no current at all. That null condition is the basis of one of the most precise measuring instruments in the elementary laboratory. The next section takes the balanced Wheatstone bridge and its practical cousin, the metre bridge, and turns the null condition into a method for measuring an unknown resistance to high accuracy.

THE WHEATSTONE BRIDGE AND THE METRE BRIDGE

Direct measurement always affects the circuit being measured. An ohmmeter sends a small test current through the unknown resistor and reads back the voltage that develops; a voltmeter taps a tiny current out of the circuit it is supposed to be observing. Both instruments are reasonably good, but neither is exact, because the current they draw alters the very quantity they are trying to read.

A **null method** avoids this fault. The instrument is wired so that, when correctly adjusted, no current flows through the detector. There is then nothing being drawn from the circuit, and the balance condition gives the unknown quantity in terms of known ones. The price you pay is a small loss of speed: you must adjust until the detector reads zero. The gain is precision better than any direct meter reading.

This section develops the two most important null-method instruments for resistance measurement. The Wheatstone bridge in its laboratory form uses four resistors and a galvanometer; the metre bridge is the same circuit rebuilt around a uniform resistance wire, which makes the balance length easy to read directly from a centimetre scale. The closing subsection treats the unbalanced bridge, showing how the galvanometer current changes when the balance condition is not met. The unbalanced case needs the full Kirchhoff machinery of the previous section.

The Wheatstone bridge

Four resistors R_1, R_2, R_3, R_4 are arranged as the four sides of a diamond, with corners A, B, C, D. An EMF source is connected across one diagonal AC, and a sensitive galvanometer is connected across the other diagonal BD. The next figure shows the layout.

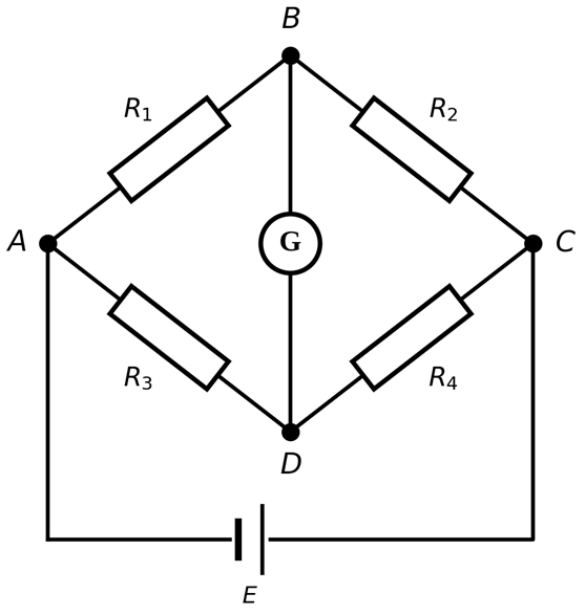


Figure: ***The Wheatstone bridge.*** *Four resistors R_1, R_2, R_3, R_4 sit on the sides of the diamond ABCD. The EMF E drives current in at A and out at C. A sensitive galvanometer G connects B to D and detects any imbalance.*

Three of the four resistors are known precisely. The fourth, the unknown, is what we are trying to measure. One of the three known resistors is a calibrated variable resistance box. Adjust this variable resistance until the galvanometer reads exactly zero. When that happens, the bridge is said to be balanced, and the unknown resistance follows from the balance condition we are about to derive.

At balance, no current flows through the galvanometer. The current must therefore flow only along the two parallel paths A → B → C through R_1 and R_2, and A → D → C through R_3 and R_4. Call these currents I_{12} (through R_1 and R_2 in series) and I_{34} (through R_3 and R_4 in series).

Zero galvanometer current also means $V_B = V_D$ (no current flows through R_g, so by Ohm's law applied to the galvanometer there is zero voltage across it). This gives us two equations. The voltage drop across R_1 (from A to B) equals the voltage drop across R_3 (from A to D):

$$I_{12}R_1 = I_{34}R_3$$

The voltage drop across R_2 (from B to C) equals the voltage drop across R_4 (from D to C):

$$I_{12}R_2 = I_{34}R_4$$

Divide the first equation by the second; the currents cancel and we are left with the balance condition:

$$\mathbf{\frac{R_1}{R_2} = \frac{R_3}{R_4}}$$

This is the working equation of the Wheatstone bridge. If three of the four resistances are known, the fourth is determined. For instance, if R_1 is the unknown:

$$R_1 = \frac{R_3}{R_4} \times R_2$$

The precision of the measurement is the precision of the three known resistors plus the sensitivity of the galvanometer to small off-balance currents. The EMF source need not be precisely known: the balance condition does not involve E. This is one reason the method is so accurate. The voltage of the driving battery may drift, the temperature of the room may change, even the galvanometer's own resistance does not enter the answer. Only the three known resistors set the precision.

The metre bridge

The Wheatstone bridge with four resistors is theoretically perfect but practically clumsy. To balance it, you have to keep adjusting a calibrated variable resistance box, which is bulky and expensive. A clever rearrangement of the same circuit replaces two of the resistors with a uniform resistance wire one metre long, and replaces the variable resistance with a sliding contact (jockey) that touches the wire at any chosen position. This is the metre bridge.

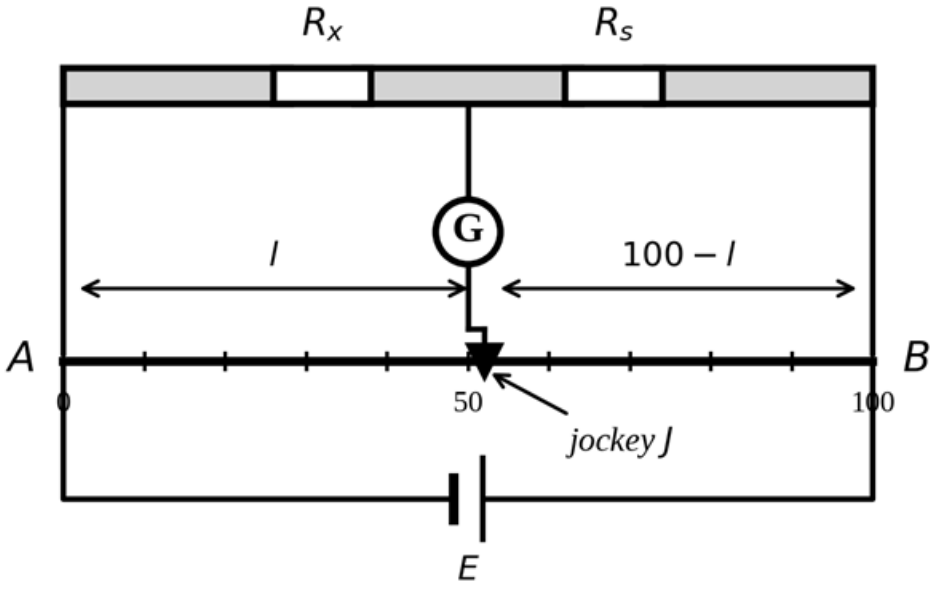


Figure: ***The metre bridge.*** *A uniform 1m resistance wire is stretched between two metallic strips A and B. The unknown R_x is clamped into one gap; a known standard resistor R_s is clamped into the other. A galvanometer G is connected between the middle strip and a sliding jockey J that can touch the wire at any position. At balance, the jockey is at distance l from A. The driving EMF E sits below the wire.*

The unknown resistance R_x sits in one gap and a known standard R_s sits in the other. The two segments of the metre wire (from A to the jockey, and from the jockey to B) play the roles of R_3 and R_4 of the four-resistor bridge.

Why a uniform wire is so useful: because the wire has uniform cross-section and uniform material, its resistance is proportional to its length.

$$R = \frac{\rho l}{A}; \text{for constant } \rho \text{ and } A, R = \text{constant} \times l, \text{and therefore:}$$

$$R \propto l$$

So the resistance of the segment from A to the jockey at position l centimetres is proportional to l, and the resistance of the segment from the jockey to B is proportional to $(100 - l)$ centimetres. The constant of proportionality (the resistance per centimetre) is the same for both segments and cancels out when we take their ratio.

Move the jockey along the wire until the galvanometer reads zero. At that position the bridge is balanced. Apply the balance condition from the four-resistor analysis:

$$\frac{R_x}{R_s} = \frac{R_l}{R_{100-l}} = \frac{l}{100-l}$$

(both l and $100 - l$ in centimetres). Solving for the unknown:

$$\mathbf{R_x = \frac{l}{100-l} \times R_s}$$

The balance length l is read directly off a centimetre scale running underneath the wire, so a single measurement of length (plus the known standard resistor) gives the unknown resistance. No variable resistance box is needed, which is the practical advantage of the metre bridge over its laboratory cousin.

Practical tip on choosing the standard: *For best accuracy, choose the standard R_s so that the balance point falls near the middle of the wire. Reading l with a metre rule has the same absolute error wherever the jockey sits (a fraction of a millimetre, say), but the fractional error in $l/(100 - l)$ is smallest when l and $100 - l$ are both large. If the unknown is small (a few ohms), use a small standard. If it is large, use a large standard. The aim is to keep the balance length between about 40 cm and 60 cm.*

The unbalanced Wheatstone bridge

What happens when the bridge is not balanced? The galvanometer no longer reads zero; it deflects by an amount that depends on the four bridge resistors, the galvanometer resistance, and the supply EMF. The deflection is not the unknown we want, but it tells us how far the bridge is from balance and which way to adjust. Finding the size of the galvanometer current in an unbalanced bridge needs the full method of the previous section, Kirchhoff's laws.

In an unbalanced bridge five currents flow: four in the bridge arms and one in the galvanometer. The five currents are linked by KCL at the two junctions B and D and by KVL around two independent loops. Solving the resulting system of equations gives all five currents, including the galvanometer current I_g of greatest interest.

The balanced bridge is a special case of this calculation: if the balance condition R1/R2 = R3/R4 holds, the Kirchhoff equations yield $I_g = 0$ automatically. The balance condition is the algebraic statement of that vanishing. Conversely, any non-zero galvanometer current is the bridge telling you it is unbalanced; the size of I_g measures by how much.

Worked Example 36 demonstrates the full Kirchhoff treatment for an unbalanced bridge with specific numerical values. The node-voltage method works particularly well here, and we shall use it.

Three examples follow.

HOT Example 34

In a metre bridge experiment, two unknown resistors R_1 and R_2 are placed in the left and right gaps respectively. The first null point is found at 25 cm from end A. A 12Ω resistor is then connected in parallel with R_2, and the null point shifts to 40 cm from A. Find R_1 and R_2.

Solution

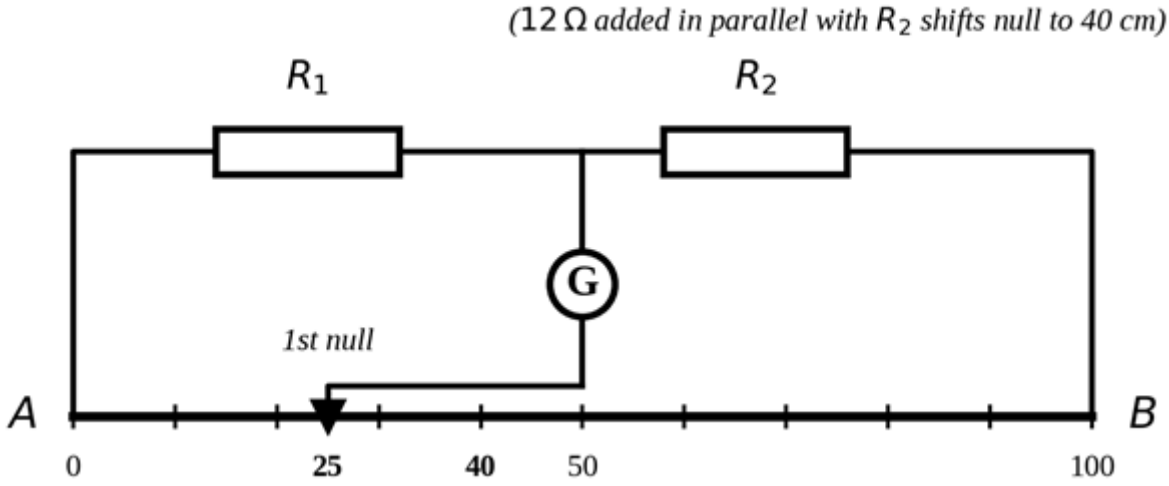


Figure: ***Example 34.*** *R_1 in the left gap and R_2 in the right gap. The first null is found at 25 cm; when a 12Ω resistor is added in parallel with R_2, the null shifts to 40 cm.*

Two balance conditions, two unknowns. Write each balance condition using the derived metre bridge formula.

First balance, null at 25 cm:

$$\frac{R_1}{R_2} = \frac{25\text{cm}}{100\text{cm} - 25\text{cm}} = \frac{25\text{cm}}{75\text{cm}} = \frac{1}{3}$$

So:

$$R_2 = 3R_1 \quad \ldots(\text{i})$$

For the second balance, a 12Ω resistor is now in parallel with R_2. The effective resistance in the right gap is the parallel combination:

$$R_{2,\text{eff}} = \frac{12\Omega \times R_2}{12\Omega + R_2}$$

The second balance, with the null now at 40 cm, reads:

$$\frac{R_1}{R_{2,\text{eff}}} = \frac{40\text{cm}}{100\text{cm} - 40\text{cm}} = \frac{40\text{cm}}{60\text{cm}} = \frac{2}{3}$$

Rearranging:

$$R_{2,\text{eff}} = \frac{3}{2}R_1 \quad \ldots(\text{ii})$$

Substitute the expression for $R_{2,\text{eff}}$ and use equation (i) to eliminate R_2:

$$\frac{12\Omega \times 3R_1}{12\Omega + 3R_1} = \frac{3}{2}R_1$$

Solve for R_1:

$$R_1 = 4\Omega$$

And from equation (i):

$$R_2 = 3 \times 4\Omega = 12\Omega$$

Checking both balances: First balance: $\frac{4\Omega}{12\Omega} = \frac{1}{3} = \frac{25}{75}$, so null at 25 cm. ✓

Second balance, with $R_{2,\text{eff}} = \frac{12\Omega \times 12\Omega}{12\Omega + 12\Omega} = 6\Omega$: $\frac{4\Omega}{6\Omega} = \frac{2}{3} = \frac{40}{60}$, so null at 40 cm. ✓

Making Sense of the Answer: *Adding a 12Ω resistor in parallel with the 12Ω right-gap resistor halves the right-gap resistance from 12Ω to 6Ω. The balance ratio $l/(100-l)$ must double to compensate, going from $1/3$ (which is $25/75$) to $2/3$ (which is $40/60$). The null point has shifted rightward from 25 cm to 40 cm, which is exactly what halving the right-gap resistance demands.*

Think Like a Physicist: *This problem is a two-stage application of the same balance equation. The trick is to use the simplest possible form of the balance condition at each stage and to write each as a fraction before doing any algebra. The fractions $1/3$ and $2/3$ carry the information of both null lengths in a compact form, and they combine cleanly into a quadratic in R_1 that has only one positive root.*

HOT Example 35

A resistor X is placed in the left gap of a metre bridge and a fixed resistor S is placed in the right gap. The resistor X is gradually heated. At $30°\text{C}$ the balance point is at 51.5cm from the left end; at $100°\text{C}$ the balance point is at 54.6cm. Find the temperature coefficient of X and the balance point if X were cooled to $0°\text{C}$.

Solution

Use the metre bridge balance equation at each temperature. At temperature θ, with balance length l:

$$\frac{X(\theta)}{S} = \frac{l}{100\text{cm} - l}$$

Substitute the two given balance lengths:

$$\frac{X(30°\text{C})}{S} = \frac{51.5\text{cm}}{48.5\text{cm}}$$

$$\frac{X(100°\text{C})}{S} = \frac{54.6\text{cm}}{45.4\text{cm}}$$

Take the ratio of the two equations; S cancels, leaving a direct ratio of resistances at the two temperatures:

$$\frac{X(100°\text{C})}{X(30°\text{C})} = \frac{54.6\text{cm} \times 48.5\text{cm}}{45.4\text{cm} \times 51.5\text{cm}} = 1.1326$$

Now apply the ratio form of the temperature dependence introduced earlier in the chapter (referenced to $0°\text{C}$):

$$\frac{X(\theta_1)}{X(\theta_2)} = \frac{1 + \alpha\theta_1}{1 + \alpha\theta_2}$$

With $\theta_1 = 100°\text{C}$ and $\theta_2 = 30°\text{C}$:

$$1.1326 = \frac{1 + 100°\text{C} \times \alpha}{1 + 30°\text{C} \times \alpha}$$

Solve for α:

$$\alpha = 2.01 \times 10^{-3}/°\text{C}$$

For the balance length at $0°\text{C}$, first compute $X(0°\text{C})/X(30°\text{C})$ using the ratio form:

$$\frac{X(0°C)}{X(30°C)} = \frac{1}{1 + 30°C \times \alpha} = \frac{1}{1 + 30°C \times 2.01 \times 10^{-3}/°C} = 0.9431$$

Then the ratio $\frac{X(0°C)}{S}$ follows from $\frac{X(30°C)}{S} = \frac{51.5cm}{48.5cm} = 1.0619$:

$$\frac{X(0°C)}{S} = \frac{X(0°C)}{X(30°C)} \times \frac{X(30°C)}{S} = 0.9431 \times 1.0619 = 1.0015$$

At balance, the metre bridge equation gives:

$$\frac{l}{100cm - l} = \frac{X(0°C)}{S} = 1.0015$$

Solving for l:

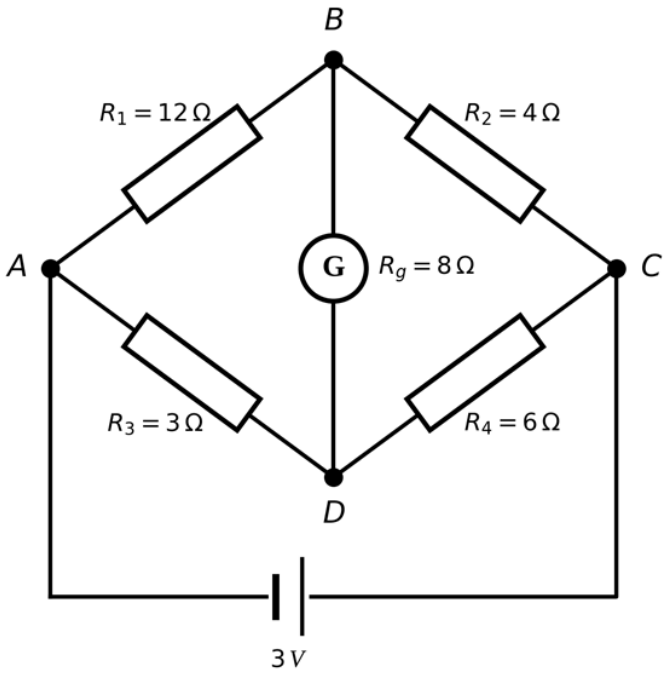
$$l = 50.04cm$$

Making Sense of the Answer: *The temperature coefficient* $2.0 \times 10^{-3}/°C$ *is in the range expected for ordinary metals (copper is* $4.3 \times 10^{-3}/°C$*, aluminium is* $4.0 \times 10^{-3}/°C$*, nichrome is* $4 \times 10^{-4}/°C$*). The value* $2 \times 10^{-3}/°C$ *sits between these, suggesting an alloy of moderate temperature sensitivity. At* $0°C$ *the balance length is very nearly 50 cm, which makes physical sense: at this temperature X and S turn out to be almost equal in resistance, and equal-resistance gaps put the null at the midpoint of the wire.*

Think Like a Physicist: *Taking the ratio of the two balance equations was the key step. It eliminated both* S *(which is just a standard reference) and the unknown* X_0 *(the cold resistance of* X*). The remaining equation had only one unknown,* α*, which then dropped out cleanly. When a problem gives you two balance conditions or two measurements at different conditions, consider their ratio as the first move; most of the constants cancel.*

HOT Example 36

A Wheatstone bridge has arms $R_1 = 12\Omega$, $R_2 = 4\Omega$, $R_3 = 3\Omega$, $R_4 = 6\Omega$, supplied by a 3V battery of negligible internal resistance. A galvanometer of resistance $R_g = 8\Omega$ is connected between B and D. Find the galvanometer current and interpret what its value indicates about the balance condition.



Solution

First check the balance condition: $R_1/R_2 = 12\Omega/4\Omega = 3$, but $R_3/R_4 = 3\Omega/6\Omega = 0.5$. The two ratios differ by a factor of 6, so the bridge is very far from balance and a substantial current will flow through the galvanometer.

Use the node-voltage form of Kirchhoff's laws. Take $V_C = 0$ as the reference, so $V_A = 3V$ (the source has no internal resistance). The unknowns are V_B and V_D. Write KCL at each.

At node B, current arrives from A through R_1 and leaves through R_2 to C and through R_g to D:

$$\frac{3 - V_B}{12} = \frac{V_B}{4} + \frac{V_B - V_D}{8}$$

At node D, current arrives from A through R_3 and (possibly) from the galvanometer, and leaves to C through R_4:

$$\frac{3 - V_D}{3} + \frac{V_B - V_D}{8} = \frac{V_D}{6}$$

Multiply the node-B equation by 24:

$$2(3\text{V} - \text{V}_\text{B}) = 6\text{V}_\text{B} + 3(\text{V}_\text{B} - \text{V}_\text{D})$$
$$11\text{V}_\text{B} - 3\text{V}_\text{D} = 6 \quad \ldots(\text{i})$$

Multiply the node-D equation by 24:

$$8(3\text{V} - \text{V}_\text{D}) + 3(\text{V}_\text{B} - \text{V}_\text{D}) = 4\text{V}_\text{D}$$
$$\text{V}_\text{B} - 5\text{V}_\text{D} = -8\text{V} \quad \ldots(\text{ii})$$

Solving with a calculator gives:

$$\text{V}_\text{D} = 1.808\text{V}$$
$$\text{V}_\text{B} = 1.040\text{V}$$

Since $\text{V}_\text{D} > \text{V}_\text{B}$, the current through the galvanometer flows from D to B (against the direction we initially assumed). Its magnitude is:

$$\text{I}_\text{g} = \frac{\text{V}_\text{D} - \text{V}_\text{B}}{\text{R}_\text{g}} = \frac{1.808\text{V} - 1.040\text{V}}{8\Omega} = 0.096\text{A} = 96\text{mA}$$

Making Sense of the Answer: *Nearly a tenth of an ampere flows through a device meant to detect microamperes. A real laboratory galvanometer would be instantly destroyed by such a current. The problem is academic rather than experimental: it shows that when the arm ratios are highly mismatched (*3 *versus* 0.5*, a factor of* 6*), the galvanometer current is large. In real Wheatstone experiments, the bridge is brought near balance first (often by protective resistors in series with the galvanometer that are shorted out as the balance is approached), and the final fine adjustment is made when I_g is already small.*

Think Like a Physicist: *The size of the galvanometer current is a quantitative measure of how far the bridge is from balance. A balanced bridge gives $I_g = 0$; a slightly unbalanced bridge gives a small I_g; a wildly unbalanced bridge gives a huge I_g. The direction of I_g tells you which way to adjust: if it flows from D to B, you need to lower R_1/R_2 or raise R_3/R_4. The bridge gives both magnitude and direction information at the same time, which is what makes it so useful as a precision comparator.*

Three examples illustrate the bridge in three modes. In every case the precision of the result rests on a length measured along a wire or a ratio of known resistances, never on the EMF of the driving cell or the sensitivity of the galvanometer. That is the strength of a null method.

The next section extends the null-method principle from resistance measurement to EMF measurement. The instrument is called a potentiometer, and it answers a question no voltmeter can answer exactly: *what is the true EMF of a cell, separately from the voltage drop across its internal resistance?*

THE POTENTIOMETER

The previous section gave us a null method for resistance. This section gives us a null method for EMF. The instrument is called the potentiometer, and it does for cells what the Wheatstone bridge does for resistors: it measures the unknown quantity without drawing any current from the source being measured. The consequence is that a potentiometer reads the true EMF of a cell, not the terminal voltage. This distinction matters whenever a cell has appreciable internal resistance, which is to say, almost always.

Recall the terminal-voltage relation derived earlier in the chapter: when a cell of EMF E and internal resistance r drives a current I through a load, the voltage across its terminals is $\text{V} = \text{E} - \text{Ir}$. A voltmeter, however good, always draws some current; it always reads V, never E. The error is small when I is small or r is small, but it is never zero. The potentiometer escapes this trap by making I exactly zero through the cell at the moment of measurement.

Why the potential gradient along the wire is uniform

A potentiometer is built around a long uniform resistance wire, stretched along a metre scale. A separate driving cell pushes a steady current along the whole length of the wire. Because the wire is uniform in cross-section and uniform in material, the current density is the same everywhere along it, and Ohm's law applied locally gives a uniform potential gradient: equal lengths of wire carry equal voltage drops.

Quantitatively, let the wire have length L, cross-section A, and resistivity ρ. Its resistance per unit length is:

$$\text{r}_\text{L} = \frac{\rho}{\text{A}}, \text{a constant} \left(\text{From R} = \frac{\rho \text{l}}{\text{A}};\ \text{r}_\text{L} = \frac{\text{R}}{\text{l}} = \frac{\rho}{\text{A}}\right)$$

With a steady driving current I_d flowing, the potential drop per unit length is:

$$\mathbf{k} = \text{I}_\text{d} \times \text{r}_\text{L} = \frac{\mathbf{I_d \rho}}{\mathbf{A}}$$

This quantity **k** is the **potential gradient** along the wire, in volts per metre or volts per centimetre depending on the unit of length. Because all three factors on the right are constants, k is constant along the wire.

Equivalently, the total potential drop across the wire is $\text{k} \times \text{L}$, which (if no other resistances are in the driving loop) equals the **EMF of the driving cell $\mathbf{E_d}$**. So:

$$E_d = k \times L \text{ or } k = \frac{E_d}{L}$$

Once the potential gradient is known, the potential at any point along the wire (measured from end A) is k times the distance. Touching a jockey at distance l from A picks up a potential of $k \times l$. This is the working principle of every potentiometer measurement: balance an unknown voltage against $k \times l$ for some length l, and the unknown voltage is determined by the balance length.

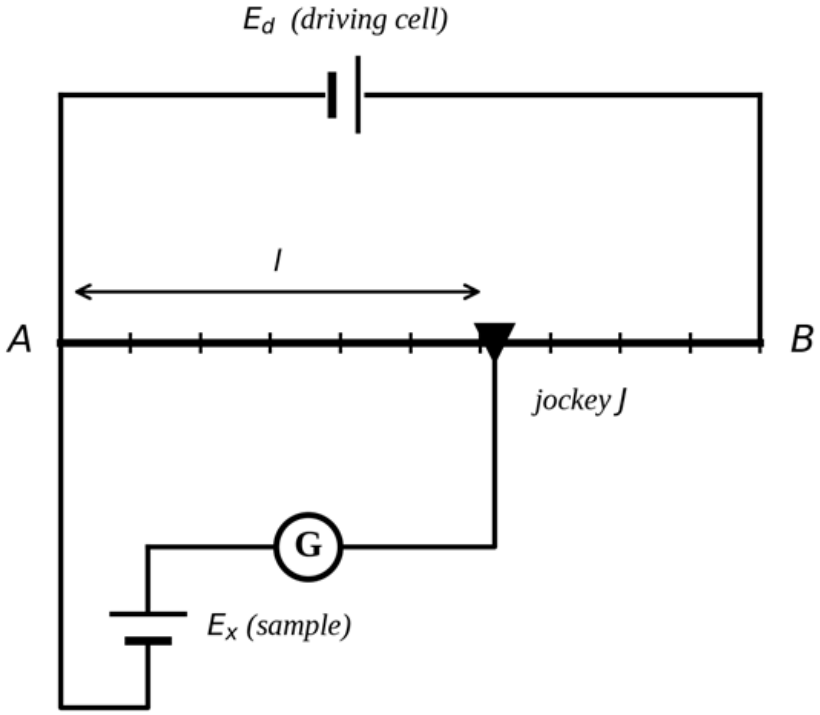


Figure: ***The basic potentiometer.*** *A driving cell E_d pushes a steady current along the uniform wire AB. The jockey J taps the wire at distance l from A. The galvanometer G connects the jockey to one terminal of the sample cell E_x; the other terminal of E_x returns to A. When the jockey is at the balance position, no current flows through G and the sample cell's EMF equals $k \times l$.*

At the balance point, the potential of the jockey (measured relative to A, along the wire) exactly equals the EMF of the sample cell. The galvanometer therefore has zero voltage across it, so by Ohm's law no current flows through it. Because no current flows through the galvanometer arm, no current flows through the sample cell either; the sample cell is instantaneously on open circuit, and its terminal voltage equals its EMF. The balance reading is therefore the true EMF, not the terminal voltage under load.

Three measurements follow from this principle: comparison of EMFs, measurement of internal resistance, and calibration of the instrument. Each is treated separately below.

Comparison of EMFs

Suppose we have a standard cell of accurately known EMF Es, and an unknown cell of EMF Ex. We can connect them in turn to the potentiometer (using a two-position switch) and find the balance length for each in turn. Let Ls be the balance length for the standard cell, and Lx be the balance length for the unknown cell.

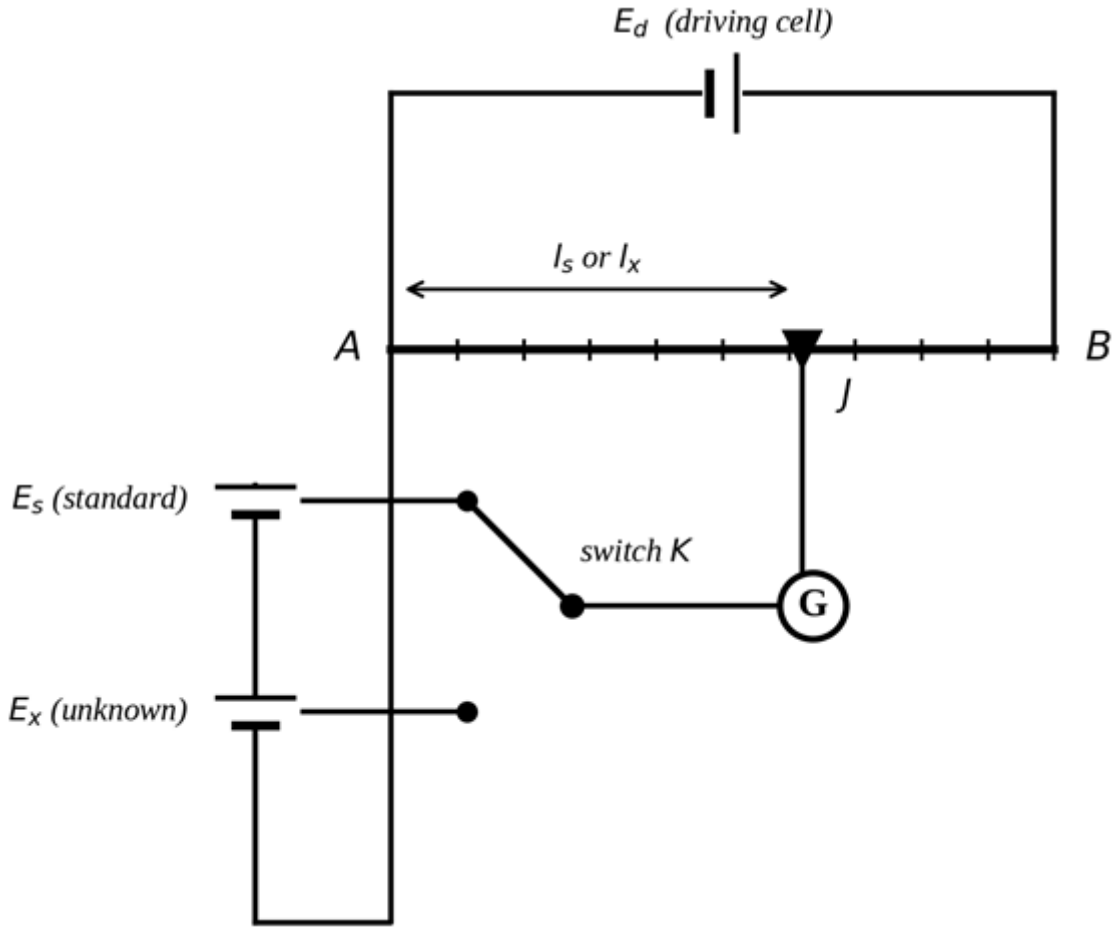


Figure: ***Potentiometer arrangement for comparing two EMFs.*** *A two-position switch K selects either the standard cell E_s or the unknown cell E_x; the chosen cell is balanced on the wire. The balance length for each is measured.*

At balance:

$$E_s = k \times L_s \text{ and } E_x = k \times L_x.$$

Dividing the second by the first cancels k:

$$\frac{E_x}{E_s} = \frac{L_x}{L_s}$$

Rearranging for the unknown EMF:

$$\mathbf{E_x = \frac{L_x}{L_s} \times E_s}$$

The unknown EMF is determined entirely by the two balance lengths and the standard EMF. The driving current I_d, the resistance per unit length, and any other property of the wire do not enter; they all cancelled with k. The same is true of the galvanometer resistance and the switch resistance, neither of which carries any current at balance.

Measurement of internal resistance

The internal resistance of a cell cannot be measured directly with an ohmmeter (the cell's own EMF would interfere with the test current). The potentiometer offers an indirect route. Take two balance readings of the same cell: one on open circuit (no external load), and one with a known external resistor R connected across the cell.

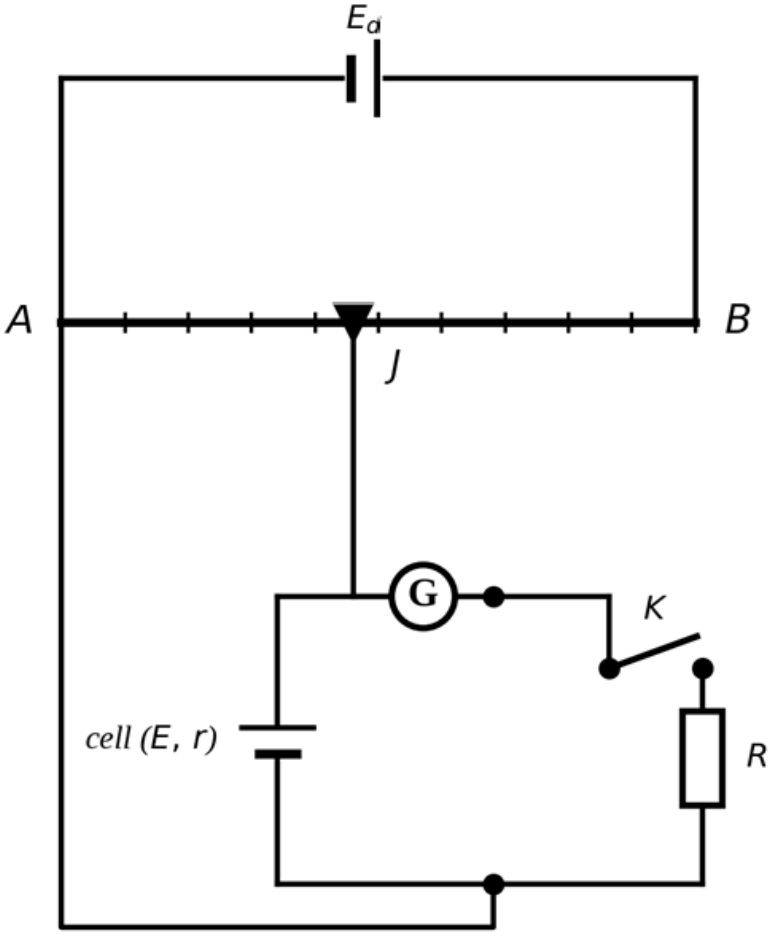


Figure: ***Potentiometer arrangement for measuring the internal resistance of a cell.*** *With the switch K open, the cell is on open circuit and its EMF is balanced at length L. With K closed, the resistor R draws current from the cell, the terminal voltage drops below the EMF, and the new balance length L1 is smaller. The internal resistance follows from the two lengths and R.*

Open-circuit reading first. With the switch K open, the cell delivers no current (the galvanometer arm carries none at balance, and the load R is disconnected). The terminal voltage therefore equals the EMF E. Let L be the balance length under this condition. Then:

$$E = k \times L$$

Now close the switch K. Current flows from the cell through R (this is the cell's load current). The cell's terminal voltage drops below the EMF by the internal-drop Ir:

$$V = E - Ir$$

The current flowing through R is $I = \frac{V}{R}$, so:

$$V = E - \frac{V}{R} \times r$$

Multiply through by R and rearrange:

$$VR = ER - Vr$$
$$V(R + r) = ER$$
$$\frac{E}{V} = \frac{R + r}{R}$$

Re-balance the potentiometer with the load now connected. The new balance length L_1 balances the terminal voltage V (the galvanometer arm carries no current at balance, so the voltage across the cell's terminals at balance equals the potential of the jockey, which is $k \times L_1$). So:

$$V = k \times L_1$$

Taking the ratio E/V using $E = k \times L$ and $V = k \times L_1$:

$$\frac{E}{V} = \frac{L}{L_1} = \frac{R + r}{R}$$

Solve for r:

$$r = \frac{RL}{L_1} - R = R\left(\frac{L}{L_1} - 1\right)$$

$$\mathbf{r = R\left(\frac{L - L_1}{L_1}\right)}$$

The internal resistance is determined by the two balance lengths and the known load R. The driving current and the potential gradient k cancel out of the answer, just as in the EMF comparison.

Calibration of the potentiometer

Before the potentiometer can be trusted to measure an unknown EMF directly (rather than by comparison), the potential gradient k must be known. Calibration consists of balancing a standard cell of known EMF E_s at length L_s and computing:

$$k = \frac{E_s}{L_s}$$

Once k is known, any unknown EMF balanced at length L_x follows from:

$$E_x = k \times L_x$$

Verification: A second standard cell of different EMF balanced at a different length should give the same k. If it does not, the driving cell has drifted, the wire is not uniform, or there is a loose contact in the driving loop. A potentiometer that fails this check should not be trusted until it passes again.

Three examples now follow.

BINDER Example 37

A potentiometer is first balanced with a standard cell of EMF $E_s = 1.50\text{V}$ at length $L_s = 40\text{cm}$. The same potentiometer then balances a second cell A at length $L_A = 60\text{cm}$. Find the EMF of cell A, and explain briefly why the potentiometer ensures an accurate comparison.

Solution

Apply the EMF-comparison formula derived in this section:

$$\frac{E_A}{E_s} = \frac{L_A}{L_s}$$

Substitute the given numerical values:

$$E_A = \frac{L_A}{L_s} \times E_s = \frac{60\text{cm}}{40\text{cm}} \times 1.50\text{V} = 2.25\text{V}$$

Explanation: The potentiometer ensures an accurate comparison because at balance no current flows from cell A through the galvanometer arm, so cell A is momentarily on open circuit and the potential picked off the wire equals the true EMF of cell A, not the terminal voltage it would show while delivering current to a load.

Making Sense of the Answer: *Cell A has a larger EMF than the standard cell, so it requires a longer length of wire to balance. The ratio of balance lengths,* $\frac{60cm}{40cm} = 1.5$*, is exactly the ratio of the two EMFs, which is why* $E_A = 1.5 \times 1.50V = 2.25V$*. The number is reasonable: 2.25V is a little more than a single dry cell's nominal 1.5V, the sort of value a fresh cell or a small two-cell source might give.*

Think Like a Physicist: *The whole comparison rests on a single fact: the potential along a uniform wire grows in direct proportion to the distance from the start. So a ratio of two EMFs becomes a ratio of two lengths, and the potential gradient* k *cancels without ever needing to be measured. A voltmeter could give the same answer in principle, but only by drawing a small current and so reading each cell's terminal voltage rather than its EMF. Whenever you need a true EMF and not a loaded voltage, reach for a balance method, not a meter.*

HOT Example 38

In measuring the internal resistance of a cell with a potentiometer, the EMF is first balanced by a length of 90cm on the potentiometer wire (with the cell on open circuit). When a 5Ω resistor is connected across the cell's terminals, the terminal voltage now balances at 45cm. (a) Explain why the balance length is smaller when the resistor is connected. (b) Calculate the internal resistance of the cell.

Solution

(a) When the resistor R is connected, the cell drives a current I through R, and the terminal voltage of the cell drops below its EMF by the internal-drop Ir. A smaller voltage requires a shorter length of wire to balance.

(b) Apply the internal-resistance formula:

$$r = R \times \frac{L - L_1}{L_1},$$

with $L = 90\text{cm}$, $L_1 = 45\text{cm}$, and $R = 5\Omega$:

$$r = 5\Omega \times \frac{90\text{cm} - 45\text{cm}}{45\text{cm}} = 5\Omega \times \frac{45\text{cm}}{45\text{cm}} = 5\Omega$$

Making Sense of the Answer: *The balance length halved when the load was connected, dropping from $90cm$ to $45cm$. This means the terminal voltage under load is exactly half the EMF: $V = E/2$. From $V = E - Ir$ with $V = E/2$, we get $E - Ir = E/2$, so $Ir = E/2 = V$. The internal drop equals the load drop, which means the internal resistance equals the external load: $r = R = 5\Omega$. This is a special case worth recognising:* ***whenever the balance length halves on adding a load, the internal resistance equals the load resistance, regardless of the actual EMF value.***

Think Like a Physicist: *Notice how everything in the answer cancelled except R and the ratio of lengths. The EMF E never had to be known, the potential gradient k never had to be known, and the driving current I_d never had to be known. This is what* ***null methods*** *give you: answers that depend only on the known quantities the experimenter actually controls, with all the inconvenient ones cancelling out.*

HOT Example 39

A potentiometer is calibrated using a standard cell of EMF 1.0V; the balance length is 40cm. An unknown dry cell balances at 30cm under open-circuit conditions. The dry cell is then connected to an external resistor $\text{R} = 4\Omega$; the new balance length under load is $\text{L}_\text{t} = 20\text{cm}$. Find (a) the potential gradient k in V/cm, (b) the EMF of the dry cell, (c) the terminal voltage of the dry cell under load, (d) the current through R, and (e) the internal resistance of the dry cell.

Solution

(a) The potential gradient follows from the calibration. The standard cell's EMF balances at 40cm, so the potential drop per centimetre of wire is:

$$\text{k} = \frac{\text{E}_\text{s}}{\text{L}_\text{s}} = \frac{1.0\text{V}}{40\text{cm}} = 0.025\text{V/cm}$$

(b) The EMF of the dry cell is the open-circuit balance reading. At $\text{L}_\text{x} = 30\text{cm}$:

$$\text{E}_\text{x} = \text{k} \times \text{L}_\text{x} = 0.025\text{V/cm} \times 30\text{cm} = 0.75\text{V}$$

(c) The terminal voltage under load balances at $\text{L}_\text{t} = 20\text{cm}$:

$$\text{V} = \text{k} \times \text{L}_\text{t} = 0.025\text{V/cm} \times 20\text{cm} = 0.50\text{V}$$

(d) The current through R follows from Ohm's law applied to the load:

$$\text{I} = \frac{\text{V}}{\text{R}} = \frac{0.50\text{V}}{4\Omega} = 0.125\text{A}$$

(e) The internal resistance follows from the terminal-voltage relation $\text{V} = \text{E} - \text{Ir}$, which rearranges to:

$$\text{r} = \frac{\text{E}_\text{x} - \text{V}}{\text{I}} = \frac{0.75\text{V} - 0.50\text{V}}{0.125\text{A}} = 2.0\Omega$$

Making Sense of the Answer: *The five quantities form a tightly linked story. The driving wire carries a potential gradient of 0.025V/cm, which converts every length into a voltage. The unknown cell has EMF 0.75V (open circuit, 30cm) and terminal voltage 0.50V (under 4Ω load, 20cm). The drop from 0.75V to 0.50V is the internal-drop $\text{Ir} = 0.25\text{V}$, and with the current $\text{I} = 0.125\text{A}$ we recover $\text{r} = 2.0\Omega$. The internal resistance is half the external load, which means two-thirds of the cell's EMF appears across the load and one-third is lost as the internal drop. A fresh dry cell would have a smaller internal resistance; 2Ω suggests this one is approaching the end of its useful life.*

Think Like a Physicist: *This problem strings together five separate physical quantities, but every one is derived from a length measurement on the potentiometer wire and an Ohm's-law step. The potentiometer has done two distinct things in one experiment: measured the EMF (by open-circuit balance) and measured the terminal voltage under load (by closed-circuit balance). Their difference, divided by the load current, gives the internal resistance. The whole experiment uses only lengths, a known standard EMF, and one known load resistor; no voltmeter or ammeter is needed anywhere.*

The potentiometer turns three separate measurements (EMF, terminal voltage, internal resistance) into a single experiment based on balance lengths along a wire. The accuracy of every answer rests on the same two pillars: a uniform wire (giving a uniform potential gradient) and a null detection (giving zero current from the cell being measured at the moment of measurement). When both conditions hold, the answers are independent of the driving cell, the galvanometer, and every property of the apparatus except the lengths and the standard reference.

The next section turns from measurement to the energy budget of an electrical circuit: where the power comes from, where it goes, and how much of it can be made to do useful work.

ENERGY, POWER, MAXIMUM POWER TRANSFER, AND JOULE HEATING

Every section so far has treated the circuit as a place where charge moves and voltage falls. This section treats it as a place where energy is spent, which, as it happens, is the only thing your electricity bill has ever cared about. A source pours energy into the circuit at some rate; the circuit hands that energy out to its components at the same rate; and in a resistor the energy handed out does not pile up or get stored, it turns at once into heat. The questions that matter from here on are accounting

questions: *how fast the source supplies energy, how fast each component spends it, what fraction of the supplied energy reaches the load the circuit was built to serve, and what is lost on the way.*

Those four questions decide why a kettle element glows, why a fuse melts at exactly the right moment, why an amplifier runs hot enough to warm your hands, and why TANESCO sends power across the country at 132kV instead of the 240V that reach a house. All of it follows from one idea: a resistor is a device that turns electrical energy into heat, and it does so at a rate we can write down exactly.

Where the heat comes from

Imagine a crowd of people shoving their way down a long corridor crammed with furniture. They do not glide through; they bump, they jostle, they lose a little energy at every collision, and the furniture and the air grow warmer for it. That corridor is a resistor, the crowd is the conduction electrons, and the warmth left behind is the entire subject of this section. Nothing in the picture changes when we make it quantitative. We only count the energy that the crowd leaves behind.

So count it. A resistor R carries a steady current I, with a voltage V across its ends. In a time t, a quantity of charge $Q = It$ passes through it. Each coulomb of that charge falls through the potential difference V as it crosses the resistor, and falling through V joules per coulomb means each coulomb gives up V joules. The total energy given up is the charge multiplied by the voltage drop:

$$\mathbf{U = QV}$$

Now substitute the two facts we already know about a resistor: the charge that flowed is $Q = It$, and the voltage across it obeys Ohm's law, $V = IR$. Putting both into $U = QV$:

$$\mathbf{U} = QV = (It)(IR) = \mathbf{I^2Rt}$$

This is the energy delivered to the resistor in time t, and because a resistor stores nothing, it is also the heat produced in the resistor in that time. *The conversion of electrical energy into heat inside a current-carrying resistor* is called **Joule heating**, and the quantity I^2Rt is the **Joule heat**.

The earlier warm lamp wire was Joule heating seen from the outside; this is the same effect written as a formula. Every drifting electron that collides with the lattice surrenders a little of the energy the field gave it, the lattice takes that energy up as vibration, and the wire grows warm. The formula I^2Rt simply counts the total of all those tiny surrenders over the whole time t.

The three forms of electrical power

Energy is the whole story, but rate is what we usually feel. A heater that delivers a megajoule over a week keeps no one warm; the same megajoule in ten minutes is a serious fire. So before any formula, fix the idea: ***power is the rate of spending energy**, and a resistor spends it as heat.*

Dividing the Joule heat by the time gives that rate, the power dissipated in the resistor:

$$P = \frac{U}{t} = I^2R$$

This is the first of three forms. Because the resistor also obeys $V = IR$, the same power can be written two other ways, depending on which pair of quantities you happen to know. Replacing one factor of I by $\frac{V}{R}$, or replacing IR by V, gives the full set:

$$P = I^2R = VI = \frac{V^2}{R}$$

These are not three different powers. They are one power wearing three outfits, and you pick the outfit that matches the data. If a problem hands you current and resistance, use I^2R. If it hands you voltage and current, use VI. If it hands you voltage and resistance, use $\frac{V^2}{R}$. Choosing the form that fits the data saves a step of algebra and avoids inventing a quantity the problem never gave you. The unit of power is the watt, one joule per second, and from $P = VI$ one watt is also one volt-ampere.

Two of the forms carry a warning worth stating once. In $P = I^2R$ the power rises with the **square** of the current, so doubling the current through a fixed resistor quadruples the heat; a small rise in current can cook a wire that felt cool a moment earlier. In $P = \frac{V^2}{R}$ the power rises with the square of the voltage but **falls** as the resistance rises, so for a fixed supply voltage a smaller resistance draws more power, not less. The kettle element with the lower resistance is the more powerful one, a fact that feels backwards right up until you decide to trust the formula over your gut.

The energy budget of a source

A resistor only consumes. A source supplies, and it is not quite as generous as it looks: it skims a little off the top to warm itself. To see the whole budget, return to the cell of EMF E and internal resistance r driving a current I through an external load R, the circuit met earlier in the chapter. The EMF is the work the source does per coulomb, so the rate at which the source delivers energy to the entire circuit, the total power generated, is:

$$P_{source} = EI$$

That total splits in two, because the current must fight its way through the internal resistance before it ever reaches the load. The power wasted inside the source, heating the cell itself, is the Joule heating in r:

$$P_{internal} = I^2 r$$

What is left over reaches the external load. Subtracting, and using $E - Ir = V$, the terminal voltage:

$$\mathbf{P_{external}} = P_{source} - P_{internal} = EI - I^2 r = I(E - Ir) = \mathbf{VI}$$

So the power reaching the load is the terminal voltage times the current, exactly the VI you would compute at the load's own terminals, with no reference to what goes on inside the source. The budget closes: every joule the EMF supplies either heats the internal resistance or reaches the load, and nothing goes missing.

The fraction that reaches the load is the **efficiency** of the transfer, *the external power divided by the total power generated:*

$$\eta = \frac{P_{external}}{P_{source}} = \frac{VI}{EI} = \frac{V}{E} = \frac{R}{R + r}$$

The last step uses $V = IR$ and $E = I(R + r)$, which lets the current cancel and leaves efficiency as a pure ratio of resistances. Efficiency is high when R is large compared with r, because then most of the EMF lands on the load and little is lost inside. It sinks toward one half, and below, as R shrinks toward r. This one ratio, $\frac{R}{R+r}$, is the hinge on which the rest of the section turns.

When is the most power delivered to the load?

Efficiency is one question; raw power is another, and they do not share an answer. Imagine the source is fixed, its EMF E and internal resistance r beyond your control, and the only knob you may turn is the load R. You want the load to receive as much power as it possibly can. Turn R too low and the load is nearly a short circuit with almost no voltage across it; turn R too high and you choke the current to a trickle. Somewhere between greed and restraint there is a best setting, and we can find it exactly.

Write the load power in terms of R alone. The current is $I = \frac{E}{R+r}$, and the power in the load is I^2R:

$$P = I^2 R = \frac{E^2 R}{(R + r)^2}$$

As R climbs from zero, this power starts at zero, rises to a peak, and falls back toward zero. To pin the peak, differentiate P with respect to R and set the derivative to zero. Using the quotient rule on $\frac{E^2 R}{(R+r)^2}$:

$$\frac{dP}{dR} = \frac{E^2[(R + r)^2 - R \cdot 2(R + r)]}{(R + r)^4}$$

The numerator simplifies. Cancel one factor of $(R + r)$ from top and bottom, and collect the bracket $(R + r) - 2R = r - R$:

$$\frac{dP}{dR} = \frac{E^2(r - R)}{(R + r)^3}$$

The denominator is always positive, so the derivative vanishes only when its numerator does, that is when $r - R = 0$:

$$R = r$$

The load receives the most power when its resistance equals the internal resistance of the source. This is the **maximum power transfer theorem**. The derivative is positive for $R < r$ (power still climbing) and negative for $R > r$ (power now falling), which confirms the turning point is a maximum, not a minimum. Substituting $R = r$ back into the power gives the height of the peak:

$$\mathbf{P_{max}} = \frac{E^2 r}{(r + r)^2} = \frac{E^2 r}{4r^2} = \mathbf{\frac{E^2}{4r}}$$

And the efficiency at the matched point follows from the ratio above, with $R = r$:

$$\eta = \frac{R}{R + r} = \frac{r}{2r} = \frac{1}{2} = 50\%$$

Here is the sting in the tail. At the very setting that pours the most power into the load, exactly half of the source's total power is being burned inside the source itself. Maximum power and high efficiency are not the same goal; they sit at opposite ends of one curve. The matched load grabs the largest possible slice, but it pays by forcing the source to waste an equal slice internally. To run efficiently instead, you must make the load far larger than r, accepting less power in exchange for wasting less.

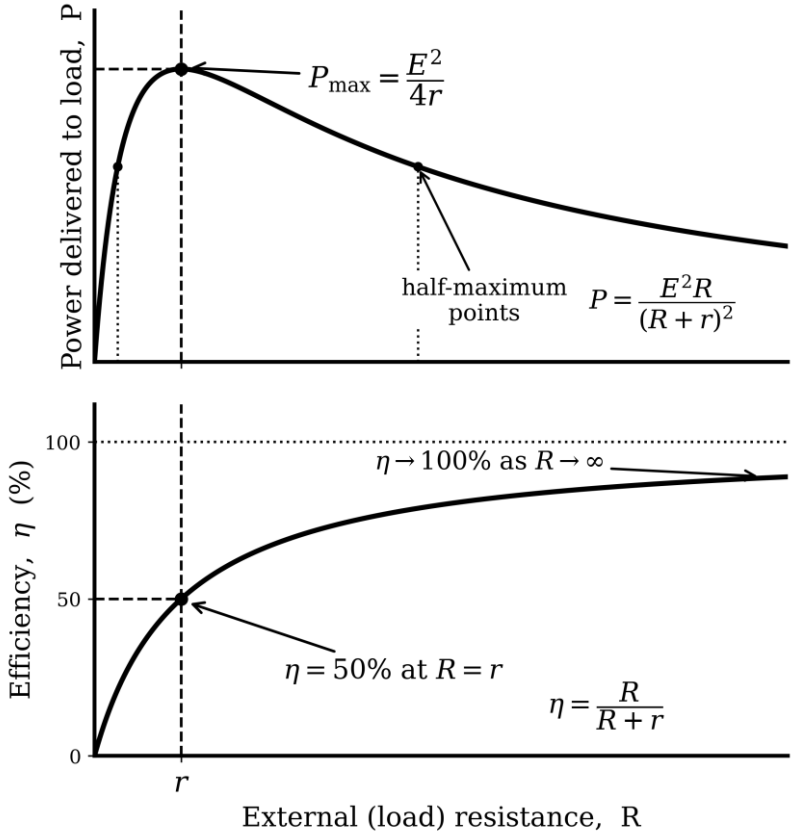


Figure: ***The power delivered to the load and the efficiency of the transfer, both plotted against the load resistance R, with the internal resistance r marked on the axis.*** *The power (upper curve) climbs from zero to a single peak directly above $R = r$ at the height $\frac{E^2}{4r}$, then falls away slowly; its half-maximum points are marked for context. The efficiency (lower curve) rises steadily, crosses exactly 50% at $R = r$, and climbs toward 100% as R grows without bound. At the one value $R = r$ the power is greatest while the efficiency is only one half.*

Kipute saw the contradiction at once.

Kipute: *Sir, if the maximum-power condition wastes half the source's energy as heat inside the source, why is anyone ever interested in maximum power transfer? Fifty per cent efficiency sounds like a failure, not a target.*

Mr. Akilikubwa: *It is a failure in a power station, and a triumph in a loudspeaker. The trick is to ask what is scarce. In an audio amplifier the source's power is cheap and plentiful; what is precious is the power pushed into the speaker, because that is what becomes sound. So the amplifier's output is deliberately matched to the speaker's four or eight ohms, accepting fifty per cent efficiency to win the loudest possible note. The wasted half is exactly why the back of an amplifier is warm.*

Kipute: *And TANESCO does the opposite.*

Mr. Akilikubwa: *It must. The source there is the whole Mtera plant, and wasting half of it would be a national disaster, so the grid runs at well over ninety-five per cent efficiency, with the load resistance enormously larger than the line's. Same theorem, opposite design. The maximum power theorem tells you the trade-off; it never tells you which side of it to stand on. That choice is engineering, not physics.*

Designing things that get hot on purpose

Joule heating is a nuisance in a transmission line and the entire point of a heater. A heating appliance is nothing more than a resistor chosen so that I^2R, or equivalently $\frac{V^2}{R}$, comes out to the wattage the appliance is meant to deliver at the voltage it is meant to run on. Designing one is choosing a resistance, then choosing a wire that has that resistance, runs hot without melting or burning away, and barely changes its resistance as it heats.

That last requirement quietly rules out copper, the very metal used everywhere else in the circuit. Copper's resistivity is low, so a copper element would need an absurd length of wire to reach a useful resistance, and copper's resistance climbs steeply with temperature, so it would draw one power cold and a different one hot, and oxidise away at red heat.

The standard answer is **nichrome**, an alloy of **ni**ckel and **chrom**ium whose resistivity is roughly sixty times that of copper, so a short coil suffices, and whose temperature coefficient of resistance is small, so its power barely drifts from cold to

glowing. Nichrome also grows a protective oxide skin that lets it sit at bright red heat in open air for years. Every exposed heating coil you have ever seen, in a room heater, a hair dryer, a toaster, is nichrome for these three reasons at once.

Around that core idea sit the familiar appliances.

- **An electric kettle** hides a sealed element in the water and is rated by the power it must deliver, from which its resistance follows by $R = \frac{V^2}{P}$; a 2.5kW kettle on 240V mains is a resistance of about 23Ω.
- **An electric iron** uses the same kind of element but adds a thermostat, a bimetallic strip that bends as it warms and breaks the circuit at a set temperature, then remakes it as it cools, holding the soleplate at the chosen heat by switching the Joule heating on and off.
- **A fuse** is the same physics turned to protection: a short length of low-melting-point wire, often a lead-tin alloy, sized so that at the rated current its I^2R heating keeps it just below melting, but at any larger current the extra heat melts it and breaks the circuit before the fault can set the rest of the wiring alight.
- **An immersion heater** is a long nichrome wire wound on a ceramic former and sheathed in a metal tube, sized by exactly the same $\frac{V^2}{R}$ sum that sizes the kettle.

Why power is sent across the country at high voltage

The last application is the largest of all. A power station sits far from the towns it feeds, and the line that carries its power has resistance, perhaps a hundred ohms over a hundred kilometres of conductor. That line resistance sits in series with everything, so the current that feeds the load must also pass through the line, heating it by I^2R. This heat is pure waste: it warms the open countryside and serves no one. The whole art of transmission is to make that waste small.

Here is the key, and it is worth slowing down for. The same power can be carried either as a small current at a high voltage or a large current at a low voltage, because power is the product VI. But the line loss depends only on the current, not on the voltage. Write the loss as I^2R_{line} and replace the current by $I = \frac{P_{load}}{V_{load}}$:

$$\mathbf{P_{line}} = I^2R_{line} = \left(\frac{\mathbf{P_{load}}}{\mathbf{V_{load}}}\right)^2 \times \mathbf{R_{line}}$$

Read this slowly, because it contains the entire reason for the national grid. *For a fixed power delivered and a fixed line, the loss falls as the **square** of the transmission voltage.* Double the voltage and the loss drops to a quarter; raise it tenfold and the loss drops to a hundredth.

To carry a fixed power, a higher voltage means a proportionally smaller current, and since the loss goes as the current squared, a high voltage with a small current is dramatically cheaper to transmit than a low voltage with a large current. This is why the grid steps the voltage up to hundreds of kilovolts for the long haul, then steps it back down with transformers to the safe 240V at the point of use.

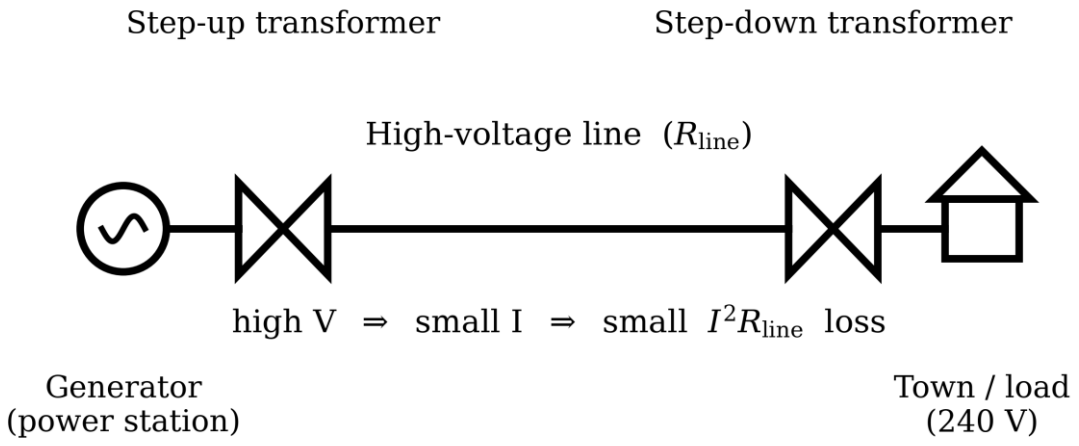


Figure: ***A high-voltage transmission line drawn from left to right.*** *A generator at the power station feeds a step-up transformer that raises the voltage to a high transmission value. A long line of resistance R_{line} carries the power across country at high voltage and therefore small current, keeping the I^2R_{line} loss small. A step-down transformer at the far end lowers the voltage again, delivering it to the town at $240V$.*

The theory is now complete, which means it is time to discover whether it survives contact with actual numbers. Six worked examples follow. Work each one before you read the solution; physics, like swimming, is not learned by watching other people do it. The solutions will wait for you. They are extremely patient, being made entirely of ink.

BINDER Example 40

An electric lamp is marked $100\text{W}, 240\text{V}$. It is switched on for one hour. Find (a) the current it draws, (b) the resistance of its filament at working temperature, and (c) the total electrical energy it converts to heat and light in the hour, in joules and in kilowatt-hours.

Solution

(a) The marking tells us the lamp dissipates 100W on 240V. Power, voltage and current are tied by $\text{P} = \text{VI}$, so:

$$\text{I} = \frac{\text{P}}{\text{V}} = \frac{100\text{W}}{240\text{V}} = 0.417\text{A}$$

(b) The resistance follows from the form of the power equation that uses the two quantities we were handed, voltage and power:

$$\text{R} = \frac{\text{V}^2}{\text{P}} = \frac{(240\text{V})^2}{100\text{W}} = 576\Omega$$

(c) Energy is power times time. One hour is 3600s:

$$\text{U} = \text{Pt} = 100\text{W} \times 3600\text{s} = 360000\text{J} = 360\text{kJ}$$

In the unit your bill actually uses, one kilowatt-hour is one kilowatt sustained for one hour, so a 100W lamp burning for one hour uses 0.1kWh.

Making Sense of the Answer: *The three parts are three readings of one lamp. It draws less than half an ampere, which is why the thin flex to a reading lamp never feels more than warm. Its working resistance of 576Ω is far above the cold resistance a meter would show, because tungsten resistance climbs steeply with temperature and the filament runs well past two thousand degrees. And the $360kJ$ it spends in an hour is mostly heat, not light: an incandescent lamp is a heater that happens to glow, which is the whole reason it has been retired in favour of cooler technologies.*

Think Like a Physicist: *Notice how the choice of formula followed the data. Part (a) had power and voltage and wanted current, so $P = VI$ was the natural form. Part (b) had voltage and power and wanted resistance, so $\frac{V^2}{R}$ was natural. At no point was it necessary to find the current first and then the resistance from it, though that route would also work. Reading the data before reaching for a formula is half the skill; the three forms of the power equation exist precisely so that whatever pair you are handed, one of them fits without a detour.*

HOT Example 41

A battery of EMF 12V has an internal resistance of 0.5Ω. (a) Find the power delivered to an external resistor of 4Ω, and the efficiency of the transfer. (b) Repeat for an external resistor of 0.5Ω. (c) Show that 0.5Ω is the load that draws the maximum power, and state that maximum. (d) Comment on the efficiency at the matched load compared with part (a).

Solution

(a) With $\text{R} = 4\Omega$ and $\text{r} = 0.5\Omega$, the current is:

$$\text{I} = \frac{\text{E}}{\text{R} + \text{r}} = \frac{12\text{V}}{4\Omega + 0.5\Omega} = 2.67\text{A}$$

The power in the load is I^2R, and the efficiency is $\frac{\text{R}}{\text{R+r}}$:

$$\text{P} = \text{I}^2\text{R} = (2.67\text{A})^2 \times 4\Omega = 28.4\text{W}$$

$$\eta = \frac{\text{R}}{\text{R} + \text{r}} = \frac{4\Omega}{4\Omega + 0.5\Omega} = 0.889 = 88.9\%$$

(b) With $\text{R} = 0.5\Omega$ the current is larger, because the total resistance is smaller:

$$\text{I} = \frac{12\text{V}}{0.5\Omega + 0.5\Omega} = 12\text{A}$$

$$\text{P} = \text{I}^2\text{R} = (12\text{A})^2 \times 0.5\Omega = 72\text{W} \qquad \eta = \frac{0.5\Omega}{0.5\Omega + 0.5\Omega} = 0.50 = 50\%$$

(c) The maximum power transfer theorem says the load draws the most power when $\text{R} = \text{r}$. Here $\text{r} = 0.5\Omega$, so the matched load is precisely the 0.5Ω resistor of part (b). The maximum power is:

$$\text{P}_{\text{max}} = \frac{\text{E}^2}{4\text{r}} = \frac{(12\text{V})^2}{4 \times 0.5\Omega} = 72\text{W},$$

which agrees with the 72W found directly in part (b), confirming that 0.5Ω is the matched load.

(d) At the matched load the efficiency is only 50%: half the battery's power is burned inside the battery. The 4Ω load of part (a) delivers less power, 28.4W against 72W, but it does so at 88.9% efficiency, wasting only a ninth of the battery's energy internally. The matched load wins on power and loses on efficiency, exactly the trade-off the theorem describes.

Making Sense of the Answer: *The two loads tell the two halves of the story in numbers. The 4Ω load is the efficient regime: it is eight times the internal resistance, so most of the EMF lands on the load and only one part in nine is lost inside. The 0.5Ω load is the matched regime: it equals the internal resistance, so the EMF splits evenly and the battery heats up as much as the load does. The matched load draws more than twice the power but throws away half of everything to do it. Which one is right depends entirely on whether the battery is feeding a power-hungry device whose supply is cheap, or a circuit where battery life is precious.*

Think Like a Physicist: *When a problem mentions both maximum power and a specific internal resistance, the matched load is handed to you for free: it is simply $R = r$, with no calculation. The maximum power itself is then $\frac{E^2}{4r}$ without ever finding the current. Recognising the matched condition by sight, rather than re-deriving it each time, turns part (c) into a single substitution. The derivation is there to be understood once; the result is there to be used on sight ever after.*

HOT Example 42

A resistor wound from copper wire forms part of the control circuitry in a Tanzanian textile factory. Measured cold, in melting ice at $0°\text{C}$, its resistance is 10Ω; in boiling water at $100°\text{C}$ its resistance is 14Ω. In operation the resistor is connected across a 12V supply and is found to dissipate 6W. (a) Find the temperature coefficient of resistance of the copper. (b) Find the resistance under operating conditions. (c) Estimate the operating temperature, and comment on whether the resistor is safe.

Solution

(a) The resistance of a metal rises with temperature as $\text{R}_\theta = \text{R}_0(1 + \alpha\theta)$, where R_0 is the resistance at $0°\text{C}$ and θ is the temperature in degrees Celsius.
Using the boiling-water reading, with $\text{R}_0 = 10\Omega$ and $\text{R}_{100} = 14\Omega$ at $\theta = 100°\text{C}$:

$$14\Omega = 10\Omega \times (1 + \alpha \times 100°\text{C})$$
$$1.4 = 1 + 100\alpha \;\Rightarrow\; \alpha = 4 \times 10^{-3}°\text{C}^{-1}$$

(b) Under operation the resistor dissipates 6W across 12V. The resistance follows from the form that uses voltage and power:

$$\text{R} = \frac{\text{V}^2}{\text{P}} = \frac{(12\text{V})^2}{6\text{W}} = 24\Omega$$

(c) Feeding $\text{R} = 24\Omega$ back into $\text{R}_\theta = \text{R}_0(1 + \alpha\theta)$, with $\text{R}_0 = 10\Omega$ and $\alpha = 4 \times 10^{-3}°\text{C}^{-1}$:

$$24\Omega = 10\Omega \times (1 + 0.004\,\theta)$$
$$\theta = 350°\text{C}$$

A copper component sitting at $350°\text{C}$ is in serious trouble. Copper oxidises rapidly in air well below that, and $350°\text{C}$ is far above the temperature any insulation around the wire could survive. The resistor is not safe; it is running close to destruction, and the design that put a bare copper resistor where it must dissipate 6W has badly underestimated the heat.

Making Sense of the Answer: *The number tells a cautionary tale. The resistance has more than doubled from its ice-point value, from 10Ω to 24Ω, and for copper a doubling of resistance means a temperature rise of hundreds of degrees, because the coefficient is small. The same small coefficient that makes copper a stable conductor at room temperature makes a large resistance rise the signature of a dangerously large temperature rise. Six watts sounds trivial, the dim glow of an indicator lamp, but concentrated in a small copper coil with nowhere to shed its heat it drives the temperature to a level that would melt solder and char insulation.*

Think Like a Physicist: *This problem chains two separate pieces of the chapter: the temperature dependence of resistance from the earlier section, and the power dissipation of this one. The bridge between them is the operating resistance, found from the power data and then read backwards through the temperature law to recover the temperature. Whenever a problem gives you both a power-and-voltage operating condition and a resistance-temperature characteristic, that operating resistance is the hinge that joins them. Find it from the power side, then turn it into a temperature on the resistance side.*

REAL Example 43

A small heater for a college laboratory in Bagamoyo is to deliver 8W when connected across a 24V supply. It is to be wound from nichrome wire of cross-sectional area $2 \times 10^{-7}\text{m}^2$ and resistivity $1.1 \times 10^{-6}\Omega\text{m}$. (a) Find the resistance the element must have. (b) Find the length of nichrome wire required. (c) A technician decides to use wire of half the diameter instead, keeping the same power and voltage. Find the new length, and explain why it is shorter even though the wire is thinner.

Solution

(a) The element must dissipate 8W at 24V, so its resistance follows from voltage and power:

$$\text{R} = \frac{\text{V}^2}{\text{P}} = \frac{(24\text{V})^2}{8\text{W}} = 72\Omega$$

(b) The resistance of a wire is $\text{R} = \frac{\rho\text{l}}{\text{A}}$, so the length is $\text{l} = \frac{\text{RA}}{\rho}$:

$$\text{l} = \frac{\text{RA}}{\rho} = \frac{72\Omega \times 2 \times 10^{-7}\text{m}^2}{1.1 \times 10^{-6}\Omega\text{m}} = 13.1\text{m}$$

(c) Halving the diameter quarters the cross-sectional area, because area depends on the square of the diameter; the new area is $\frac{A}{4} = 0.5 \times 10^{-7}\text{m}^2$. The power and voltage are unchanged, so the required resistance is still 72Ω. The new length is:

$$l = \frac{RA}{\rho} = \frac{72\Omega \times 0.5 \times 10^{-7}\text{m}^2}{1.1 \times 10^{-6}\Omega\text{m}} = 3.27\text{m}$$

The thinner wire needs only a quarter of the length, 3.27m against 13.1m. This seems backwards until the formula is read carefully. The design fixes the resistance at 72Ω. Thinning the wire to a quarter of its area would, on its own, quadruple the resistance; to bring it back down to 72Ω, the length must be cut to a quarter. A thinner wire reaches the same resistance in a shorter run.

Making Sense of the Answer: *The result overturns a common guess. People expect a thinner wire to need more length, picturing thin wire as more resistive per metre, which is true: the thinner wire has four times the resistance per metre. But the design fixes the total resistance, not the resistance per metre, and four times the resistance per metre means a quarter of the metres to reach the same total. Both elements are* 72Ω*; the thin one packs it into* 3.27m*, the thick one spreads it over* 13.1m*. For a coil that must fit inside a small heater, the thinner wire is the more compact choice, which is one reason heating elements are wound from fine wire.*

Think Like a Physicist: *The safe way through a problem like part (c) is to hold fixed what the design holds fixed, and let the formula tell you the rest. Here the power and voltage are fixed, so the resistance is fixed at* 72Ω*; only then do you ask what length the new area demands. The error to avoid is reasoning about resistance per metre, which is not what the design controls. Anchor on the quantity that is actually constrained, the total resistance, and the counter-intuitive answer collapses into a single line of arithmetic.*

REAL Example 44

An immersion heater at a Tanga substation rest-house is rated 2.5kW and runs on the 240V mains. (a) Find the current it draws and the resistance of its element. (b) Find the energy it delivers to the water in thirty minutes, in joules and in kilowatt-hours. (c) Explain why the element is wound from nichrome rather than copper.

Solution

(a) The current follows from $P = VI$, and the resistance from $\frac{V^2}{P}$:

$$I = \frac{P}{V} = \frac{2500\text{W}}{240\text{V}} = 10.4\text{A}$$

$$R = \frac{V^2}{P} = \frac{(240\text{V})^2}{2500\text{W}} = 23\Omega$$

(b) Energy is power times time. Thirty minutes is 1800s:

$$U = Pt = 2500\text{W} \times 1800\text{s} = 4500000\text{J} = 4.5\text{MJ}$$

In the billing unit, 2.5kW running for half an hour is:

$$U = 2.5\text{kW} \times 0.5\text{h} = 1.25\text{kWh}$$

(c) Nichrome is chosen for three reasons copper cannot meet at once. Its resistivity is about sixty times that of copper, so the 23Ω element is a short, compact coil rather than an impractically long wire. Its temperature coefficient of resistance is small, so the element draws nearly the same 2.5kW whether it starts cold or runs at full heat, instead of drifting as it warms. And it grows a protective oxide layer that lets it sit at high temperature in hot water and steam for years without corroding away. Copper fails all three: it would need an enormous length, its resistance and power would swing wildly from cold to hot, and it would oxidise quickly.

Making Sense of the Answer: *The figures are the everyday arithmetic of a power bill. Ten and a half amperes is a heavy household current, the reason immersion heaters need their own thick cable and a dedicated switch rather than sharing a thin lighting circuit. And* $1.25kWh$ *for half an hour of heating is, at any realistic tariff, a noticeable cost, which is why heating water is one of the larger lines on a domestic bill. The* $4.5MJ$ *and the* $1.25kWh$ *are the same energy in two units: the joule for physics, the kilowatt-hour for the meter on the wall.*

Think Like a Physicist: *Energy problems reward keeping the units in plain sight. Working in joules means seconds, so thirty minutes had to become* $1800s$ *before multiplying by watts. Working in kilowatt-hours means hours and kilowatts, so the same calculation in those units is the cleaner one for a bill: kilowatts times hours, with no large powers of ten to carry. Knowing both routes, and which unit the question wants, is part of reading the problem correctly.*

REAL Example 45

TANESCO must deliver 1MW of power to a town 100km from a generating station. The transmission line has a total resistance of 100Ω. (a) If the power were sent at the household voltage of 240V, find the current in the line and the power lost in it, and comment. (b) If instead the power is sent at 132kV, find the current, the line loss, and the fraction of the transmitted power that is lost. (c) State the principle the comparison illustrates.

Solution

(a) At 240V, the current needed to carry 1MW is:

$$\mathrm{I} = \frac{\mathrm{P}}{\mathrm{V}} = \frac{1 \times 10^6 \mathrm{W}}{240\mathrm{V}} = 4167\mathrm{A}$$

The power dissipated in the 100Ω line at that current is:

$$\mathrm{P_{line}} = \mathrm{I^2R} = (4167\mathrm{A})^2 \times 100\Omega = 1.7 \times 10^9 \mathrm{W}$$

Comment: This is practically impossible. The line would have to dissipate 1700MW of heat to deliver 1MW to the town, more than a thousand times the power being delivered. No conductor could survive it and no station could supply it. Sending bulk power at 240V is not merely wasteful; it cannot be done at all.

(b) At 132kV, the current carrying the same 1MW is far smaller:

$$\mathrm{I} = \frac{\mathrm{P}}{\mathrm{V}} = \frac{1 \times 10^6 \mathrm{W}}{132000\mathrm{V}} = 7.6\mathrm{A}$$

$$\mathrm{P_{line}} = \mathrm{I^2R} = (7.6\mathrm{A})^2 \times 100\Omega = 5.8 \times 10^3 \mathrm{W} = 5.8\mathrm{kW}$$

The fraction of the transmitted power lost in the line is:

$$\frac{\mathrm{P_{line}}}{\mathrm{P}} = \frac{5800\mathrm{W}}{1 \times 10^6 \mathrm{W}} = 0.0058 = 0.58\%$$

Less than one per cent is lost on the way. Same line, same distance, same power delivered, but now the transmission is entirely practical.

(c) The comparison illustrates that line loss falls as the **square** of the transmission voltage for a fixed delivered power, because raising the voltage lowers the current in proportion and the loss goes as the current squared. High-voltage transmission is not a refinement or an economy; it is the only way bulk power can cross a long line at all. This is why TANESCO transmits at 132kV and above, and why every long line begins and ends at a transformer.

Making Sense of the Answer: *The two parts are separated by a factor that looks impossible until the squares are counted. The voltage rose from* $240V$ *to* $132kV$*, a factor of* 550*. The current fell by the same factor of* 550*, from* $4167A$ *to* $7.6A$*. But the loss, going as the current squared, fell by* 550 *squared, more than three hundred thousand times, from* $1.7 \times 10^9 W$ *to* $5.8 \times 10^3 W$*. A factor of* 550 *in voltage buys a factor of* 300000 *in reduced loss. That savage leverage of the square is the single fact that makes a national grid possible, and it is the same square that appeared in* $P = I^2R$ *at the very start of the section.*

Think Like a Physicist: *When a transmission problem gives you a delivered power and asks about line loss, do not start from the line resistance and a current you have not found yet. Start from* $P_{line} = \left(\frac{P}{V}\right)^2 R_{line}$*, which puts the loss directly in terms of the two things the problem fixes, the power delivered and the voltage of transmission, with the line resistance a constant. The current is an intermediate you can compute if asked, but the shape of the answer, loss falling as the inverse square of voltage, is clearest when the current is never made the hero of the calculation.*

Six problems down. If your answers matched ours, congratulations: you can now, in principle, bill an entire town for electricity. If they did not, the gap between your number and ours is precisely where the learning is hiding, so go back and hunt it down before reading on. Either way, the resistor has now surrendered every joule it intends to surrender.

CONDUCTION OF ELECTRICITY IN GASES

Every section until now has run its current through a metal, and a metal is a generous host: it comes with its conduction electrons already in place, a standing crowd waiting only for a field to tell them which way to drift. Air is not generous. Switch on a voltage across an ordinary air gap and, almost always, nothing happens at all. The air sits there as a perfect insulator, which is lucky, because it is the air between the wires in your house that stops every socket from short-circuiting through the room. Yet the same air, pushed hard enough, becomes the most violent conductor in nature and throws a lightning bolt across a kilometre of sky. This section is about that switch from insulator to conductor: why a gas refuses to carry current, what it takes to change its mind, and what the current looks like once it flows.

The empty queue

Think back to the picture from the start of the chapter, where a current was a queue of charge carriers passing buckets of energy down a line. In a metal the queue is already full. Every atom has donated an electron or two to a common pool, so the carriers are standing shoulder to shoulder before you apply any voltage; close the switch and the bucket is passed along the line at once.

A gas is the opposite. Its molecules are electrically neutral, each one holding tightly to all of its own electrons, so the queue is simply empty. There are no free charges to push, and a voltage with nothing to push moves no current. Before a gas can conduct, the carriers have to be created, and creating them means tearing electrons loose from neutral molecules. That act is called ionisation, it leaves behind a positive ion and a freed electron, and it does not happen for free: every electron is bound

to its molecule, and prying one loose costs energy. The whole of this section follows from that single sentence. *A gas conducts only when something supplies the energy to fill its empty queue with ions.*

The breakdown potential

A gas, then, will not conduct until it is ionised, and the most direct way to ionise it is to turn up the voltage until the electric field itself does the tearing. Imagine the field as a hand pulling on the molecules' electrons. At low voltage the pull is gentle and the electrons stay put. Raise the voltage and somewhere there is a threshold where the pull wins, an electron is ripped free, and that electron, accelerated by the same field, slams into the next molecule and frees another. The gap suddenly fills with carriers and the gas, an insulator a volt ago, becomes a conductor.

The voltage at which this begins is the **breakdown potential** of the gap. It can be defined as *the minimum potential difference that must be applied across a gas to cause electrical breakdown, allowing the gas to become conducting and a sustained electric current to flow through it.*

Below the breakdown potential, the gas is an insulator; above it the gas conducts, sometimes quietly in a lamp and sometimes catastrophically in a spark. The value depends on:

- the nature of the gas,
- the spacing (distance) between the electrodes,
- the shape and condition of the electrodes,
- the gas pressure, and
- the temperature of the gas.

For two rounded electrodes about 1cm apart in air it takes roughly 30kV to strike a spark. Sharpen those electrodes to points at the same separation and the figure falls to about 12kV, because the field crowds together at a sharp point and reaches the breakdown value there while the average field is still modest. The same effect sends lightning to the tip of a tall tree rather than the flat ground beside it. One more lever matters more than any of these: lowering the pressure lowers the breakdown potential sharply, which is the entire reason the discharge tubes we are about to meet are run **not** at the pressure of the room but at a small fraction of it.

What does the ionising

Turning up the voltage is only one way to fill the empty queue. Anything that can deliver enough energy to a molecule to knock an electron off it will serve as an **ionising agent**, and they fall into three families.

1) **Nuclear radiation:** The alpha, beta, and gamma rays streaming from a radioactive source each carry far more energy than the dozen or so electron-volts it takes to ionise a gas molecule, so a single ray ionises a long trail of molecules as it passes.
2) **Electromagnetic radiation from the high-energy end of the spectrum:** X-rays, carrying from about 10^2 to 10^6eV per photon, and the more energetic ultraviolet, carrying from a few up to about 10^2eV, both ionise by the photoelectric effect, the photon spending its energy to eject an electron.
3) **Accelerated charged particles:** A free electron or proton that has fallen through a potential difference of only a few tens of volts picks up enough kinetic energy that, when it collides with a neutral molecule, it ionises it on impact. This third family is the important one, because it is the seed of the avalanche: *the very electrons that one ionising event sets free are themselves accelerated by the field into ionising agents for the next collision.*

Inside a discharge tube

The cleanest place to watch a gas conduct is a **discharge tube**: a glass tube with a metal electrode sealed into each end, the cathode and the anode, wired to a high-voltage supply, with a side tube to a vacuum pump so the pressure inside can be lowered step by step. At atmospheric pressure nothing is seen. But as the pump draws the pressure down, the gas lights up, and the pattern of light marches through a remarkable sequence of stages that is worth watching closely, because each stage is a snapshot of the same physics at a different density of molecules.

At about 20mmHg the discharge first appears as ragged brushes of light at the electrodes and faint violet streamers reaching between them.

Pump down to about 5mmHg and the tube organises itself: a long pink **positive column** fills most of the length and ends in a glow at the anode, a blue **negative glow** sits near the cathode, and between them lies a dark gap called the **Faraday dark space**.

Lower the pressure again to about 0.1mmHg and the positive column retreats toward the anode and breaks into a row of bright bands called **striations**; the dark spaces stretch out, and a new dark region, the **Crookes dark space**, opens up around the cathode.

Finally, near 0.01mmHg, the positive column and the negative glow have vanished altogether, the Crookes dark space has grown to fill the whole tube, and the glass itself fluoresces with a soft green light where invisible rays from the cathode strike it.

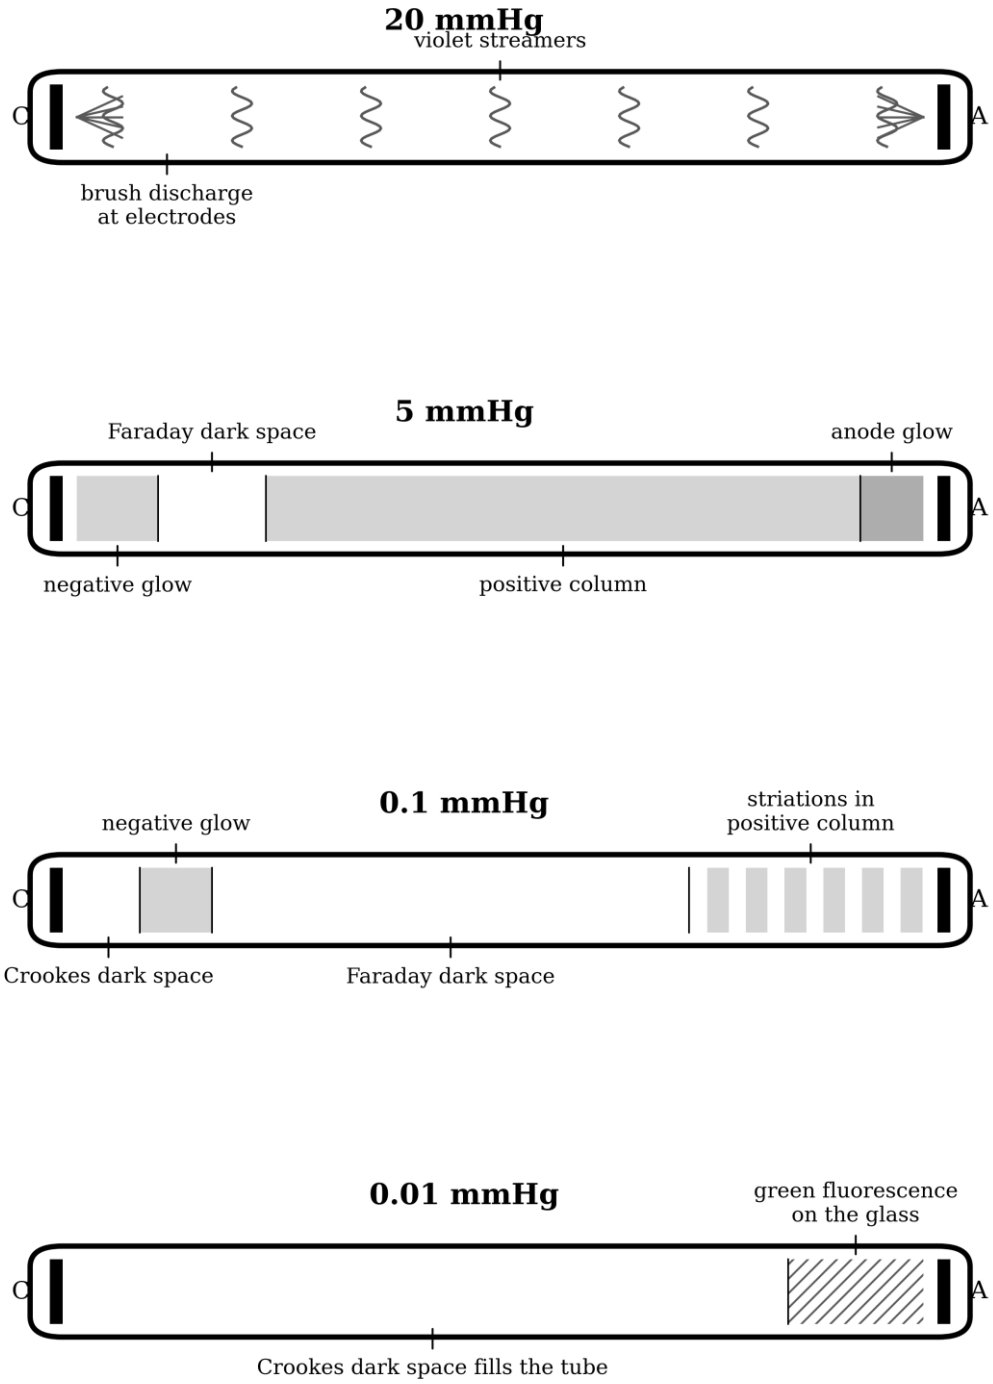


Figure: ***The appearance of the discharge in a tube as the pressure is lowered in four stages.*** *At* $20 mmHg$ *the discharge is brush-like at the electrodes with violet streamers between. At* $5 mmHg$ *a long positive column ends in an anode glow, with a negative glow near the cathode and a Faraday dark space between. At* $0.1 mmHg$ *the positive column breaks into striations and the Crookes dark space appears at the cathode. At* $0.01 mmHg$ *the Crookes dark space fills the tube and the glass walls fluoresce green. The cathode (C) is at the left and the anode (A) at the right of each tube.*

There is one experimental fact in this parade that is easy to miss and worth fixing in mind: *the length of the dark spaces depends only on the pressure, not on the length of the tube.* A longer tube does not stretch the Crookes dark space; it only lengthens the positive column to fill the extra room. The dark spaces are set by how far an electron travels between collisions, which is a matter of how crowded the gas is, and that is why the same demonstration tube can be pumped through every stage above and show each one faithfully.

Two kinds of discharge

Underneath all that scenery, a gas discharge is running in one of two fundamentally different ways, and telling them apart is the key idea of the section.

In an **initiated discharge** *the gas conducts only because some outside ionising agent, an ultraviolet lamp, an X-ray set, a radioactive source, is steadily manufacturing ions for it.* The applied voltage does an important job, sweeping those ions to the electrodes to make a current, but it does not itself create a single new ion. The proof is brutal and simple: switch off the external agent and the current dies within an instant, because the existing ions reach the electrodes or recombine and nothing replaces them. The gas is a conductor only on borrowed carriers.

In a **self-sustaining discharge** *the voltage has been raised so high that it takes over the manufacturing itself.* Now an electron accelerated by the field gains, in the short hop between one collision and the next, enough kinetic energy to ionise the molecule it strikes. That collision frees a second electron; the field accelerates both; each ionises another; and the numbers double and redouble in a runaway called **chain ionisation**. This is a positive-feedback avalanche, and once it is running the gas makes all its own carriers. You can switch the external agent off entirely and the discharge carries on, feeding itself.

The difference between the two regimes is exactly the difference between a fire you must keep lighting with a match and a fire that has caught and now spreads on its own.

The distinction is not academic. A Geiger-Müller counter is deliberately run in the self-sustaining regime, so that one single ionising particle entering the tube is enough to trigger the avalanche and produce a pulse of current large enough to hear as a click. A fluorescent lamp, by contrast, runs just short of runaway, in the region where the chain barely begins, enough to keep the current flowing steadily without the avalanche tearing away.

The ionisation curve

Both regimes, and the boundary between them, show up on a single graph. Hold the intensity of an ionising agent fixed, raise the applied voltage from zero, and plot the current that flows. The result is the ionisation curve, traditionally labelled at five points from O to S.

- From O to P the voltage is low. The agent is making ions, but the field is too weak to sweep them across before many of them meet an opposite ion and recombine, so only a fraction reach the electrodes; raising the voltage rescues more of them, and the current climbs.
- From P to Q the field is now strong enough to collect ions faster than they can recombine, and the current climbs further.
- From Q to R something new happens: every single ion the agent produces is now being collected, none are left to rescue, and so raising the voltage further does nothing at all. The curve goes flat. This plateau is the **saturation current**, and its height is set not by the voltage but by the intensity of the ionising agent, a brighter source giving a higher plateau.
- Beyond R, the voltage becomes large enough for chain ionisation to ignite, the gas begins to make its own carriers, and the current turns sharply upward toward S with no sign of levelling off. To the left of R the discharge is initiated and borrowing its carriers; to the right of R it is self-sustaining and a true plasma.

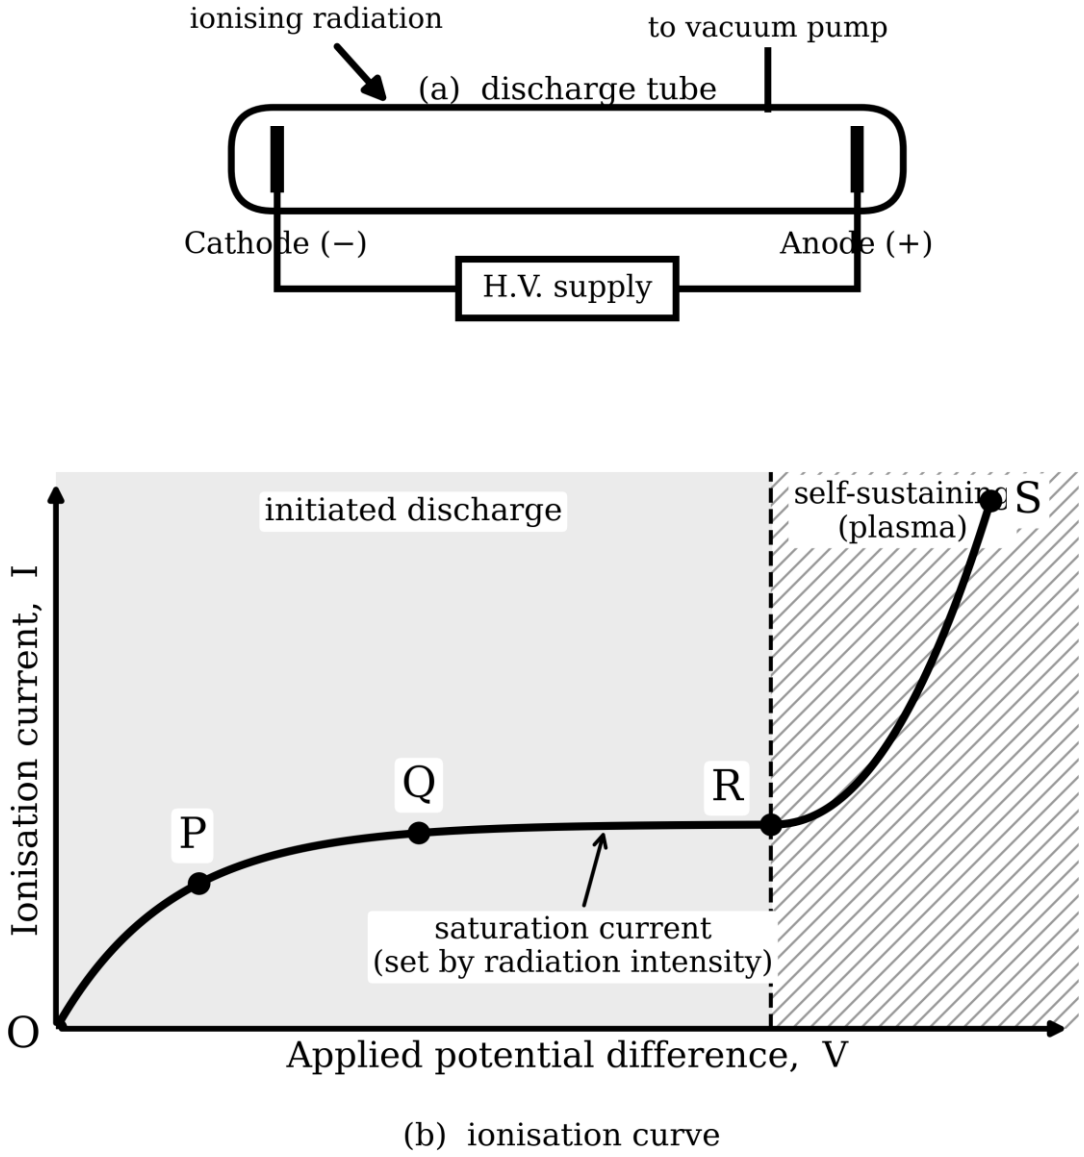


Figure: ***(a) A discharge tube, with cathode and anode wired to a high-voltage supply, a port to the vacuum pump, and an external ionising agent. (b) The ionisation curve, the current*** **I** ***plotted against the applied potential difference*** **V** ***at fixed radiation intensity.*** *From O to P to Q the current rises as fewer ions recombine; from Q to R it saturates, every ion produced now being collected, at a height fixed by the radiation intensity; beyond R, toward S, chain ionisation sets in and the current climbs steeply. The region up to R is the initiated discharge; beyond R the discharge is self-sustaining.*

The physics of this section is told in pictures and words rather than equations, so these two examples ask you to reason rather than to calculate, which is its own kind of difficulty. Read each question, look away, and try to build the answer in your own words before you check it against the provided one. A gas-discharge device you can explain out loud is a gas-discharge device you understand.

REAL Example 46

Explain why a fluorescent tube in the Miono physics laboratory flickers for a moment when it is first switched on, but then settles into a steady glow.

Solution

Two things have to come together before the tube can light, and at switch-on neither is ready. The tube is filled with mercury vapour at low pressure, and while it is cold that vapour is sparse and its breakdown potential is high, higher than the mains voltage alone can supply. Meanwhile the carriers have to be seeded: small filaments at each end of the tube are heated by the initial current and begin to emit electrons, but it takes a moment for them to warm. To bridge the gap, a starter switch in the circuit closes and then snaps open, and at the instant it opens the choke, a coil in series with the tube, responds to the sudden collapse of current with a brief high-voltage pulse. That pulse, added to the mains, finally exceeds the cold breakdown

potential and strikes the discharge. The visible flicker is exactly this struggle: the gas being driven across its threshold, perhaps stuttering through a failed attempt or two before it catches.

Once the discharge is running it warms the tube, and a warmer tube is a different device. The mercury vapour pressure rises, ionisation becomes easy, and the voltage needed to sustain the current falls well below the value needed to start it. The tube settles to a steady operating point partway along the ionisation curve, where the chain just sustains itself without running away, and the starter, no longer needed, stays open. The flicker belongs to starting; the steady glow belongs to running; and the difference between them is the difference between forcing a cold gas across its breakdown potential and coasting a warm one along its operating point.

Making Sense of the Answer: *The flicker is not a fault but a fingerprint of the physics. A cold gas has a high breakdown potential and a warm one a low operating voltage, so every fluorescent tube must briefly be over-driven to start and then runs comfortably under-driven, and the flicker is the visible seam between those two states. It is also why a tube near the end of its life flickers endlessly: its filaments no longer seed electrons well, so it never quite settles past the starting struggle.*

Think Like a Physicist: *Every gas-discharge device is a two-part problem, and naming the parts is the whole skill. One part supplies the carriers, here the heated filaments and the choke's voltage pulse; the other supplies the steady field that keeps them moving, here the mains across a now-warm vapour. When you meet any such device, sort its components into those that make carriers and those that drive them, and the operation explains itself.*

HOT Example 47

A Geiger-Müller counter operates with the gas between its electrodes held at the upper end of the R-to-S region of the ionisation curve. (a) Explain why this regime is required for the counter to register a single ionising particle. (b) Explain why the discharge extinguishes after each pulse, so that the next particle can be counted. (c) Identify the mechanism by which the discharge is quenched.

Solution

(a) A single ionising particle, on its own, frees only a handful of ions, far too few to drive a current any meter could notice. The counter therefore works in the self-sustaining region beyond R, where chain ionisation rules. Here that handful of initial ions is not the end of the story but the spark of an avalanche: each freed electron is accelerated hard enough to ionise again, the numbers double and redouble, and the single particle's faint mark is amplified into a current pulse large enough to register as an audible click. The counter does not detect the particle; it detects the avalanche the particle triggers.

(b) An avalanche that never stopped would leave the tube permanently conducting, blind to every particle after the first. So each pulse must be cut short. After the avalanche has produced its pulse, the discharge has to be extinguished, the gas returned to its non-conducting state, and the tube made ready for the next particle. Only then is the counter a counter rather than a one-shot switch.

(c) The extinguishing is done by a **quenching agent**, a small amount of a halogen or an organic vapour mixed into the tube's gas. During the avalanche, excited atoms emit ultraviolet photons, and those photons, striking the cathode or other atoms, would ignite fresh avalanches and keep the discharge alive indefinitely. The quenching vapour absorbs those ultraviolet photons before they can do so, starving the discharge of its secondary ignition and letting it die out after each pulse. Without a quenching agent the counter would latch on after a single count and never recover.

Making Sense of the Answer: *Amplification is the entire point of the device. One particle is unmeasurable and one avalanche is unmistakable, so the counter is built to convert the first into the second by living in the chain-ionisation regime. But raw amplification with no off-switch is useless, because a detector that fires once and sticks counts nothing, and the quenching gas is precisely the off-switch that turns a runaway into a repeatable pulse.*

Think Like a Physicist: *A detector of single events needs enormous gain, and enormous gain wants to run away, so every such device pairs its amplifier with a brake. Gain and reset are the two requirements, and they pull against each other exactly as raw power and efficiency did in the previous section. When you study any detector, ask the same two questions in turn: what makes the signal large, and what makes it stop, so the next signal can be seen.*

Two devices, no equations, and yet the same idea ran through both: a gas carries no current until something fills its empty queue with ions, and everything else, the flicker, the click, the colour, the quench, is a detail of how those ions are made and unmade. If you can now look at a glowing tube and say which part is making carriers and which part is driving them, the section has done its work.

The next section turns from how the gas conducts to the light it gives off while conducting. That glow is not a single colour but a precise pattern of colours, and the pattern is a fingerprint: unique to each gas, sharp enough to name an unknown sample in a laboratory, and steady enough to read the composition of a star from its light alone.

OPTICAL SPECTRA OF GASES

A discharge tube does not merely conduct; it glows, and the colour of that glow is the whole subject of this section. Neon signs burn orange-red, sodium street lamps a flat insistent yellow, mercury lamps a cold blue-white, and the colour is so dependable that a physicist can name the gas in a sealed tube from across a dark room. The previous section explained how a gas carries current; this one asks what the light it gives off while conducting can tell us. The answer is, astonishingly, almost everything about the gas, and the same trick read backwards tells us the make-up of stars no one will ever visit.

Why the glow has a colour

The first thing to clear away is a natural but wrong idea: that the colour of the glow is the colour of the gas, the way paint has a colour. It is not. The gas at rest is invisible. The colour appears only when the discharge pumps energy into the atoms, and it is the energy coming back out that we see.

Here is the mechanism, in the smallest number of moving parts. Inside every atom the electrons may occupy only certain fixed energy levels, never the gaps between them. The discharge knocks an electron up to a higher level; a moment later it falls back down, and as it falls it must shed exactly the energy it gained, which it does by emitting a single packet of light, a photon. The energy of that photon is fixed by the size of the drop, and the energy of a photon fixes its colour. A big drop gives a blue or violet photon; a small drop gives a red one. Because each element has its own private ladder of energy levels, unlike that of any other element, the set of drops it can make is unique, and so the set of colours it emits is unique. The glow is not a colour but a **pattern** of colours, and the pattern belongs to one element and no other. This is why the colour is a fingerprint: read the pattern and you have named the gas.

Emission and absorption

Spectra come in two complementary families, and the difference is simply whether the gas is giving light out or taking it in.

An **emission spectrum** is what we have been describing: a hot, excited gas radiating its own characteristic photons, seen as bright lines of colour against darkness.

An **absorption spectrum** is the mirror image. Shine white light, which contains every colour, through a cool gas, and the gas absorbs exactly the photons it would itself emit, lifting its electrons up the very same ladder. The light that comes through is then missing those colours, and the spectrum shows dark lines at precisely the wavelengths the gas would have emitted if it were hot.

The fingerprint is the same in both cases, printed in bright ink or in shadow. This mirror symmetry is the engine of astronomy: the cool outer gas of a star absorbs its own lines out of the hot light beneath, and reading those dark lines tells us which elements the star is made of, from across distances light itself takes thousands of years to cross.

Three kinds of emission spectra

Not every glowing source gives the same kind of spectrum, and the kind reveals what is doing the glowing.

A line spectrum *is a set of separate, narrow, bright lines on a dark background.* It is produced by a single-atom (monatomic) gas at low pressure, where each atom radiates alone and undisturbed, so only its own sharp transitions appear. This is the spectrum of the discharge tubes in this chapter, and the one drawn in the next figure.

A band spectrum *is a set of groups of lines, each group so closely crowded that it looks like a band.* It is produced by gases whose particles are molecules of more than one atom, such as oxygen or carbon dioxide, because a molecule can store energy not only in its electrons but in vibrations and rotations, splitting each transition into a dense cluster.

A **continuous spectrum** *is an unbroken smear of colour with no gaps at all, the full rainbow.* It is produced by hot solids, hot liquids, and gases dense enough that the atoms are packed shoulder to shoulder, their crowded energy levels blurring into one continuous span. The glowing filament of an old lamp gives a continuous spectrum; a low-pressure gas gives lines.

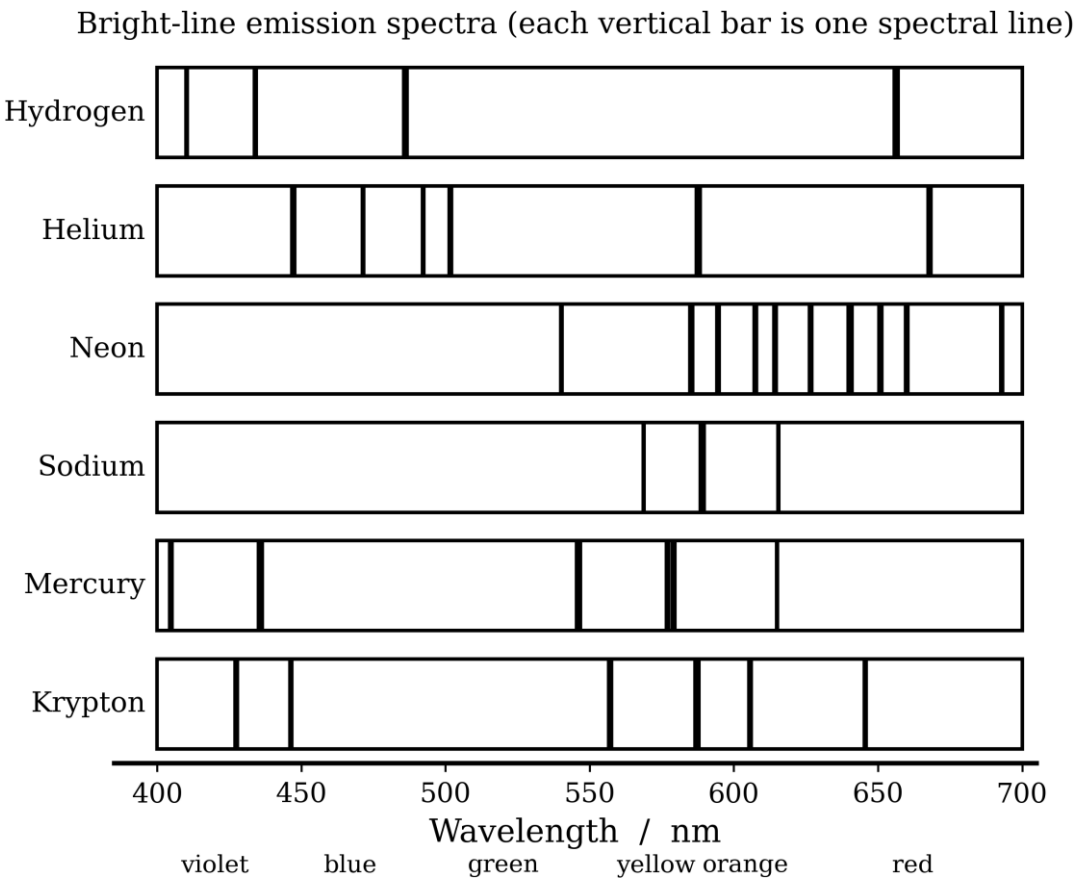


Figure: ***The bright-line emission spectra of six gases on a common wavelength axis from*** $\mathbf{400}\boldsymbol{nm}$ ***to*** $\mathbf{700}\boldsymbol{nm}$**.** *Each vertical bar is one spectral line, and the thicker bars are the stronger lines. No two patterns are alike: hydrogen shows its four Balmer lines, sodium is dominated by a single yellow line near* $589nm$*, mercury by a green line at* $546nm$*, and neon by a thicket of orange-red lines. The pattern, not any single line, identifies the gas.*

Reading the lines: the spectrometer

To turn that pattern into numbers we need an instrument that sorts light by wavelength and lets us measure each piece, and that instrument is the **spectrometer**. Before any formula, picture what it must do: take the mixed light from a discharge tube and fan it out so that each colour leaves in its own direction, then measure those directions.

The fanning out is done by a diffraction grating, a plate ruled with many thousands of fine parallel lines, met already in the chapter on waves. Light from the tube is passed through a narrow slit and a collimator that makes it into a parallel beam, the beam strikes the grating, and the grating sends each wavelength off at its own angle. A telescope mounted on a calibrated angular scale is then swung round to find each line and read the angle at which it sits. The angle and the wavelength are tied together by the grating equation from the wave chapter,

$$m\lambda = d\sin\theta$$

where d is the spacing between neighbouring rulings of the grating, θ is the angle of the line from the straight-through direction, λ is the wavelength, and m is the order, a whole number counting how many times the pattern repeats out from the centre. Measure θ for a line, and the equation hands you λ. The spectrometer turns a colour you can only describe into a wavelength you can write down.

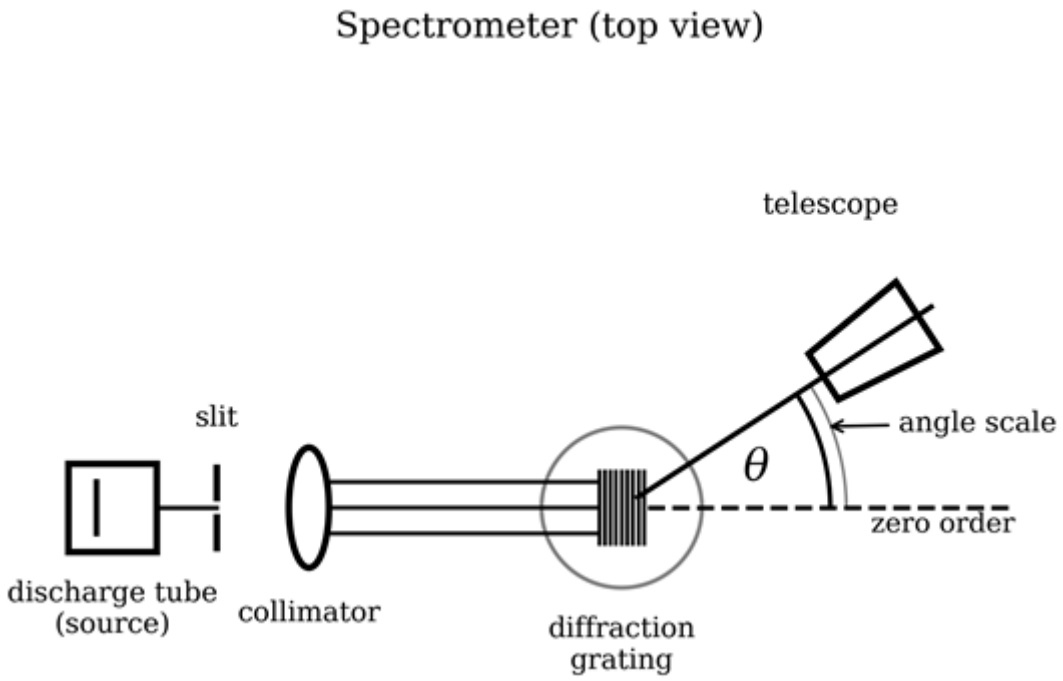


Figure: ***A spectrometer seen from above.*** *Light from the discharge tube passes through a slit and a collimator, which forms it into a parallel beam, and falls on a diffraction grating mounted on a turntable. The grating sends each wavelength off at its own angle; the straight-through direction is the zero order. A telescope on a calibrated angular scale is swung to the angle* θ *of each line, and the grating equation* $m\lambda = d\sin\theta$ *converts that angle into a wavelength.*

A table for naming the gas

Once the wavelengths are measured they are compared against known values, and identification is simply matching the pattern. The brightest visible lines of a few common gases are collected below; an unknown tube is named by measuring its lines and finding the row they fit.

Gas	**Brightest lines in the visible range: wavelength / nm (colour)**
Hydrogen	410 (violet), 434 (violet), 486 (blue-green), 656 (red)
Helium	447 (blue), 502 (green), 588 (yellow), 668 (red)
Neon	585 to 640 (a thicket of orange-red lines), 640 (red, strong)
Mercury	405 (violet), 436 (blue), 546 (green, strong), 577 and 579 (yellow)
Krypton	427 (violet), 557 (green), 587 (yellow), 606 (orange)

After a section told almost entirely in pictures, here is the one place the grating earns its keep with a calculation. Work it before you read the solution; the grating equation is short, but it has a way of hiding an extra answer or two from anyone in a hurry.

HOT Example 48

Light from a hydrogen discharge tube is examined with a diffraction grating ruled with 600 lines per millimetre. The blue-green hydrogen line, of wavelength 486nm, is one of the lines seen. (a) Find the spacing of the grating. (b) Find the angle at which the blue-green line appears in the first order. (c) Determine whether the same line can also be seen in the second order, and if so, at what angle.

Solution

(a) The grating is ruled with 600 lines in every millimetre, so the spacing between neighbouring lines is one millimetre divided by 600:

$$d = \frac{1\text{mm}}{600} = \frac{1\times10^{-3}\text{m}}{600} = 1.67\times10^{-6}\text{m} = 1.67\mu\text{m}$$

(b) The grating equation is $m\lambda = d\sin\theta$. Rearranged for the angle and used in the first order, with $m = 1$:

$$\sin\theta = \frac{m\lambda}{d} = \frac{1\times486\times10^{-9}\text{m}}{1.67\times10^{-6}\text{m}} = 0.292$$

$$\theta_1 = \sin^{-1}(0.292) = 17.0^\circ$$

(c) For the second order, set $m = 2$. The right-hand side simply doubles:

$$\sin\theta = \frac{m\lambda}{d} = \frac{2 \times 486 \times 10^{-9}\text{m}}{1.67 \times 10^{-6}\text{m}} = 0.583$$

Since 0.583 is less than 1, an angle exists and the line is indeed visible in the second order:

$$\theta_2 = \sin^{-1}(0.583) = 35.7^\circ$$

So the blue-green line appears twice, once at 17.0° and again, more widely deflected, at 35.7°.

Making Sense of the Answer: *The same line turning up at two angles is not a contradiction but the meaning of the order number m. The grating repeats its pattern outward from the centre, and each repeat throws every wavelength a little further out, so a single colour appears once in the first order and again, further from the centre, in the second. In fact this line also satisfies the equation for the third order, where $\sin\theta = 0.875$ and $\theta_3 = 61.0^\circ$; only when $\frac{m\lambda}{d}$ would exceed 1 does the line run out of orders, because no real angle has a sine greater than one.*

Think Like a Physicist: *Whenever a grating problem asks for "the angle," pause and ask "in which order?", because the honest answer is usually several angles, not one. The quantity $\frac{m\lambda}{d}$ is a sine, and a sine is capped at 1, so the test for whether an order exists is simply whether $\frac{m\lambda}{d} \leq 1$. Counting the orders this way, rather than stopping at the first, is what separates a full answer from half of one, and it is the same discipline of checking the limits of a formula that served us with maximum power and with high-voltage transmission earlier in the chapter.*

One short calculation, and the loop is closed: the colour we could only point at across a dark room has become a wavelength in nanometres, and a wavelength in nanometres has become the name of an element. The pattern that the eye reads as "orange-red" or "cold blue" is, to a spectrometer, a list of numbers as specific as a signature, and that is why the light of a flame, a lamp, or a distant star can be made to confess what it is made of.

The chapter is nearly done, and it has now built three things in turn: how charge moves through metals, how it forces its way through gases, and how the light from those glowing gases can be read. The final section spends all of it at once, walking through the everyday devices, the kettle and the fuse, the fluorescent tube and the lightning overhead, in which this chapter's physics is quietly at work.

APPLICATIONS OF CURRENT ELECTRICITY

Current electricity is at its most convincing when it stops being equations and starts being objects. This section takes the devices of everyday Tanzanian life and reads each one through the physics the chapter has built, sorting them into three families: those that work by conduction through a metal, those that work by conduction through a gas, and a few that quietly need both at once. The aim is not to add new theory but to watch the old theory earn its keep, in a kettle, a street lamp, a welder's arc, and the storm overhead.

Applications of conduction in metals

1. Electric kettle and immersion heater

The most familiar electrical device is also the simplest: a coil of nichrome or constantan, chosen through the relation $R = \frac{V^2}{P}$ so that it draws its rated power from the mains and pours that power into the water as heat at the rate $P = \frac{V^2}{R}$ until it boils. The kettle that warms a household's first cup of the day, and the immersion heater that fills the bath, are exactly the kind of element we sized earlier in the chapter; the only design decision is the resistance, and the resistance is fixed by the power and the voltage.

2. Electric iron

An electric iron is that same heated coil with one clever addition, a thermostat made from a bimetallic strip. The strip bends as it warms and opens the circuit when the soleplate reaches the chosen temperature, then closes it again as it cools, switching the Joule heating on and off so that a school uniform can be pressed at a steady heat without scorching. The element supplies the warmth; the bimetallic strip supplies the judgement.

3. Filament lamp

The oldest way of making light from electricity pushes the heating to its limit. A fine tungsten coil is run white-hot, at about 2500K, glowing simply because it is too hot not to; tungsten is chosen for its enormous melting point, above $3400\,^\circ\text{C}$, and the coil is sealed in a bulb of inert argon so that it cannot burn away in air. Its weakness is that most of its energy escapes as invisible infrared heat rather than light, which is why the filament lamp is steadily giving way across Tanzania to the cooler fluorescent and LED lamps that do the same work for a fraction of the current.

4. Fuse and circuit breaker

Not every use of heating is meant to be seen. A fuse is a deliberate weak point built into a circuit for its own protection, a short strand of low-melting lead-tin alloy sized so that its $P = I^2R$ heating keeps it just below melting at the safe current but

melts it the instant the current climbs too high, breaking the circuit before a fault can set the wiring alight; Tanzanian homes carry fuses rated at 5A, 13A, and 30A for the lighting, the sockets, and the cooker. A circuit breaker does the same duty more conveniently, tripping a switch with an electromagnet or a bimetallic strip and resetting once the fault is cleared, instead of being replaced.

5. High-voltage power transmission

On a national scale the same concern over heating decides whether power reaches the town at all. Because the loss in a line grows with the square of its current, the electricity TANESCO generates is stepped up to 132kV, and on the newer southern interconnection to 220kV, for its long journey across the country, then stepped back down at the substations to the gentle voltage that enters the house. As the chapter showed in full, low-voltage transmission over Tanzanian distances is not merely wasteful but physically impossible; the high-voltage line and its transformers are the only way bulk power can travel.

6. Multimeter shunts and multipliers

Electricity also measures itself. The multimeter on the laboratory bench reads current and voltage across wide ranges by switching extra resistors around its sensitive moving coil: a small resistor in parallel, a shunt, diverts most of the current and widens the ammeter, while a large resistor in series, a multiplier, drops most of the voltage and widens the voltmeter. Turning the range knob simply selects which shunt or multiplier is in play, exactly as the chapter developed.

7. Electroplating

A steady current can also build. In the metal workshops of every town, a direct current passed through a solution carries dissolved metal onto the surface of an object held at the cathode, plating jewellery in silver or gold and tools in chromium. The mass laid down in a given time follows Faraday's law of electrolysis, $\mathrm{m} = \frac{\mathrm{MIt}}{\mathrm{nF}}$, in which M is the molar mass, n the charge on each ion, F the Faraday constant, and It the total charge; the very charge we have tracked all chapter is here, quite literally, ferrying metal from one place to another.

Applications of conduction in gases

1. Fluorescent lighting

The second family makes light not by heating a wire but by lighting a gas, and its brightest member hangs on the wall of the laboratory itself. A fluorescent lamp is a low-pressure mercury-vapour tube whose discharge floods out invisible ultraviolet, which a phosphor painted inside the glass drinks in and returns as visible white light, far more efficiently than any glowing filament.

It is worth taking apart in the mind, because it holds nearly the whole chapter inside one glass tube, as Mr. Akilikubwa was fond of showing.

Mr. Akilikubwa (taking the spare tube down from the wall): *The lamp that lit up at the very start of the chapter will now show you the difference between metal conduction and gas conduction inside a single object.*

Kipanga: *It is only a long lamp, sir.*

Mr. Akilikubwa: *Look at the small coils sealed into its two ends. They are metal filaments, like the inside of a torch bulb. When the switch closes, current heats them, and the hot metal boils off electrons by thermionic emission, spilling free electrons into the low-pressure gas. Only now does the gas have carriers. A high-voltage pulse from the choke ionises a path between the ends, the discharge strikes, and ultraviolet light pours from the glowing mercury vapour. That light is invisible, so the glass is coated with a phosphor that swallows it and gives it back as the white you read by.*

Kipute: *So the one tube uses metal conduction in the filaments, gas conduction in the vapour, and something else again in the phosphor?*

Mr. Akilikubwa: *All three at once. The chapter has prepared you for the first two; the third, the way the phosphor turns one colour into another, waits for you in the atomic-physics chapter.*

2. Vapour discharge lamps

Strip away the phosphor and the bare colour of the gas shines through. A sodium lamp burns the hard orange of its own 589nm line over the trunk roads at night, a neon sign glows red in a Dar es Salaam shop window, and a mercury lamp throws a cold bluish-white across a factory yard. In each, the colour is nothing but the emission spectrum of the chosen gas, the fingerprint of the previous section pressed into service as signboard and street light.

3. Radiation detection: The Geiger-Müller counter

Pushed to its self-sustaining extreme, the same discharge becomes a counter rather than a lamp. The Geiger-Müller tube, met earlier, runs so near to runaway that a single ionising particle entering it triggers a whole avalanche and a current pulse loud enough to count as an audible click. This amplification of one particle into one click is why such tubes guard the staff at the Ocean Road Cancer Institute in Dar es Salaam and the radiotherapy unit at Bugando in Mwanza, wherever radioactive material is handled.

3. Plasma display panels

The same gas-discharge idea, shrunk to the size of a dot and repeated a million times across a screen, once lit every pixel of a plasma television: each cell is a microscopic discharge whose ultraviolet excites a coloured phosphor, the fluorescent lamp in miniature. Newer screens have largely replaced it, but the physics still glows in some industrial display panels.

Bridging applications: Devices that need both

A few devices put metal conduction and gas conduction to work together, in series, each doing what the other cannot.

1. Arc welding

In arc welding a fierce arc, a discharge through the open air at thousands of degrees, leaps between a metal electrode and the metal work and fuses them where it lands. The electrodes are metal conduction and the arc is gas conduction, so the welder crouched behind a dark visor at a roadside workshop is running both halves of this chapter at once.

2. Sodium street lamps with a metallic ballast

A sodium street lamp cannot run on its discharge alone. Wired in series with the glowing vapour is a metal ballast, a resistor or coil, that throttles the current the discharge would otherwise let run away and destroy itself the instant it strikes. The gas half makes the light; the metal half keeps it alive.

3. Electric arc furnace

In the electric arc furnaces of the scrap-metal yards, an arc struck between great graphite electrodes and a charge of old iron pours its heat into the metal until it runs liquid. The graphite electrodes carry the current in, and the ionised gas of the arc carries the heat, thousands of degrees of it, into the charge. Metal conduction delivers the energy; gas conduction spends it.

The physics of the chapter is now fully spent, so the last example asks the one question every household actually cares about: not how the current flows, but what it costs by the end of the month. Add the numbers yourself first, because this is the rare example you can check against the bill on your own wall.

REAL Example 49

A Tanzanian household runs the following from the 240V mains each day: a 2.5kW kettle for 12 minutes, three 9W LED bulbs for 6 hours, a 1.8kW iron for 25 minutes, and a 150W refrigerator that runs for a total of 8 hours. (a) Find the total electrical energy used each day, in kilowatt-hours. (b) At a TANESCO tariff of 290TZS for each kilowatt-hour, find the bill for a 30-day month.

Solution

(a) Energy is power times time, and if the power is in kilowatts and the time in hours, the product comes out directly in kilowatt-hours. Taking the appliances one at a time:

$$\text{E}_{\text{kettle}} = 2.5\text{kW} \times \frac{12}{60}\text{h} = 0.5\text{kWh}$$

$$\text{E}_{\text{bulbs}} = 3 \times 9\text{W} = 27\text{W} = 0.027\text{kW} \times 6\text{h} = 0.162\text{kWh}$$

$$\text{E}_{\text{iron}} = 1.8\text{kW} \times \frac{25}{60}\text{h} = 0.75\text{kWh}$$

$$\text{E}_{\text{fridge}} = 0.15\text{kW} \times 8\text{h} = 1.2\text{kWh}$$

Adding the four gives the energy used in one day:

$$\text{E}_{\text{day}} = 0.5\text{kWh} + 0.162\text{kWh} + 0.75\text{kWh} + 1.2\text{kWh} = 2.61\text{kWh}$$

(b) A 30-day month uses thirty times the daily energy, and the bill is that energy multiplied by the tariff:

$$\text{E}_{\text{month}} = 2.612\text{kWh} \times 30 = 78.4\text{kWh}$$

$$\text{Bill} = 78.4\text{kWh} \times 290\text{TZS/kWh} = 22720\text{TZS}$$

Making Sense of the Answer: *The bill is dominated by exactly the appliances that turn electricity into heat. The kettle and the iron between them, running for barely half an hour a day, account for almost half the energy, while three LED bulbs glowing all evening cost less than either, because a watt is a watt and the bulbs simply draw very few of them. The refrigerator is the quiet surprise: low-powered, but running a third of the day, it climbs to the largest single share. Energy is power multiplied by time, and time is the variable households forget.*

Think Like a Physicist: *The whole calculation is the kilowatt-hour taken seriously: kilowatts multiplied by hours, with no large powers of ten to carry, which is precisely why the electricity meter uses that unit rather than the joule. When an energy question quotes powers in watts and kilowatts and times in minutes and hours, convert everything to kilowatts and hours first, and the arithmetic becomes the simple multiplication you see here. The hardest part of a real bill is not the physics; it is remembering to count the hours honestly.*

There it is, the chapter delivered to your own front door, with a number on it in shillings. Every line of that bill is a resistor, a current, or a discharge doing exactly what the chapter said it would. Yet the story of electric charge is not confined to wires, meters, and household appliances.

Mr. Akilikubwa: *Before we close the chapter, look out of the window. A storm has gathered over the Indian Ocean to the east of Pwani, and if you watch closely you will see a bright fork of lightning streak from cloud to ground in a fraction of a second. The charge between the cloud's base and the ground has built until the electric field reached the breakdown value of air, around three million volts per metre. At that point the air ionises in a thin channel, becomes a plasma, and carries a current of tens of thousands of amperes in under a millisecond.*

From the bill on your table to the lightning in the sky, the same ideas have been at work throughout this chapter. That is the chapter's tidy account of itself, and it reads beautifully. Whether any of it has actually lodged in you is a different question entirely, and the miscellaneous worked examples that follow are where the two of us find out together.

MISCELLANEOUS WORKED EXAMPLES ON CURRENT ELECTRICITY

Example 50

(a) Explain why a copper wire conducts electricity readily whereas a glass rod does not.
(b) A conductor carries a steady current of 0.50A when a potential difference of 2.0V is maintained across it. Determine its conductance and its resistance.

Solution

(a) A metal such as copper contains a large number of free electrons that are not bound to any particular atom but are free to move throughout the lattice. When a potential difference is applied across the metal, these free electrons drift through it and so constitute a current. In glass, by contrast, every electron is tightly held within its atoms and bonds, so there are almost no free charge carriers available to move. Since a current requires mobile charge carriers, copper conducts readily whereas glass does not.
(b) Conductance is the reciprocal of resistance, and follows directly from the current and voltage:

$$\text{G} = \frac{\text{I}}{\text{V}} = \frac{0.50\text{A}}{2.0\text{V}} = 0.25\text{S}$$

The resistance is the reciprocal of the conductance:

$$\text{R} = \frac{1}{\text{G}} = \frac{1}{0.25\text{S}} = 4.0\Omega$$

The conductor has a conductance of 0.25S and a resistance of 4.0Ω.

Example 51

(a) Distinguish between an ohmic and a non-ohmic conductor, giving one example of each.
(b) A filament lamp carries a current of 0.5A when 2.0V is applied across it, and 1.0A when 8.0V is applied. Determine its resistance at each setting and account for the difference.

Solution

(a) An ohmic conductor obeys Ohm's law, so the current through it is directly proportional to the potential difference across it at constant temperature; its current-voltage graph is a straight line through the origin and its resistance stays constant. A non-ohmic conductor does not show this proportionality, so its current-voltage graph is curved and its resistance changes with the operating conditions. A metallic resistor at constant temperature is ohmic, whereas a filament lamp is non-ohmic.
(b) At the first setting the resistance follows from Ohm's law:

$$\text{R}_1 = \frac{\text{V}_1}{\text{I}_1} = \frac{2.0\text{V}}{0.5\text{A}} = 4.0\Omega$$

At the second setting:

$$\text{R}_2 = \frac{\text{V}_2}{\text{I}_2} = \frac{8.0\text{V}}{1.0\text{A}} = 8.0\Omega$$

The resistance rises from 4.0Ω to 8.0Ω because the larger voltage drives a larger current, which dissipates more power and raises the filament's temperature; and since the resistance of a metal increases with temperature, the hotter filament has the greater resistance. The lamp is therefore non-ohmic.

Example 52

(a) Explain why the resistance of a metal increases with temperature while that of a semiconductor decreases.
(b) The platinum coil of a resistance thermometer has resistance 20.0Ω at $0°\text{C}$ and 27.8Ω at $100°\text{C}$. When placed in a furnace its resistance is 51.2Ω. Determine the temperature of the furnace.

Solution

(a) In a metal the number of free electrons is essentially fixed, so conduction depends only on how freely those electrons move. As the temperature rises, the lattice ions vibrate more vigorously, the electrons collide with them more frequently, the relaxation time falls, and so the resistance increases. In a semiconductor, however, raising the temperature releases many additional charge carriers, and this large increase in the number of carriers outweighs the increased scattering, so the resistance decreases.

(b) The temperature coefficient of resistance follows from the two fixed-point readings:

$$\alpha = \frac{R_{100} - R_0}{R_0 \theta} = \frac{27.8\Omega - 20.0\Omega}{20.0\Omega \times 100°C} = 3.9 \times 10^{-3} °C^{-1}$$

Applying $R_\theta = R_0(1 + \alpha\theta)$ to the furnace reading:

$$51.2\Omega = 20.0\Omega \times (1 + 3.9 \times 10^{-3}\theta)$$

$$2.56 = 1 + 3.9 \times 10^{-3}\theta \;\Rightarrow\; \theta = \frac{1.56}{3.9 \times 10^{-3}} = 400°C$$

The furnace temperature is $400°C$.

Example 53

(a) Explain why the headlamps of a car dim momentarily when the starter motor is engaged.
(b) A 12V car battery of internal resistance 0.04Ω supplies $150A$ to the starter motor during cranking. Determine the terminal voltage during cranking and the power dissipated inside the battery.

Solution

(a) The starter motor draws a very large current from the battery. Because the battery has internal resistance, this large current produces a large voltage drop inside the battery, so the terminal voltage falls well below the EMF. The headlamps are connected across these same terminals, so the reduced terminal voltage supplies less power to them and they dim. Once the engine fires and the starter disengages, the current drops sharply, the internal voltage drop becomes small again, and the lamps return to full brightness.

(b) The terminal voltage during cranking follows from $V = E - Ir$:

$$V = E - Ir = 12V - (150A)(0.04\Omega) = 6.0V$$

The power dissipated inside the battery is the Joule heating in the internal resistance:

$$P = I^2 r = (150A)^2(0.04\Omega) = 900W$$

The terminal voltage falls to $6.0V$ during cranking, and $900W$ is dissipated as heat inside the battery.

Example 54

(a) A network containing more than one source cannot be solved by series and parallel reduction alone. Name the two laws of Kirchhoff that make such a network solvable, and identify the conservation principle that each one expresses.
(b) Two batteries are connected with their positive terminals joined together and their negative terminals joined together, so that they drive a common external resistor of 2Ω. The first has EMF $8V$ and internal resistance 1Ω; the second has EMF $7V$ and internal resistance 1Ω. Using Kirchhoff's laws, determine the current supplied by each battery and the current in the external resistor.

Solution

(a) The first is Kirchhoff's current law, which states that the sum of the currents entering any junction equals the sum of the currents leaving it; it expresses the conservation of electric charge, since charge can neither accumulate at nor be created at a junction.

The second is Kirchhoff's voltage law, which states that around any closed loop the sum of the EMFs equals the sum of the potential drops; it expresses the conservation of energy, since a unit charge taken once round a loop must return to its starting potential.

(b) Let I_1 and I_2 be the currents from the first and second batteries, and let I_3 be the current in the external resistor. Applying Kirchhoff's current law at the junction:

$$I_1 + I_2 = I_3 \ldots (i)$$

Applying Kirchhoff's voltage law to the loop containing the first battery and the resistor, with $r_1 = 1\Omega$ and $R = 2\Omega$:

$$E_1 = I_1 r_1 + I_3 R \;\Rightarrow\; 8V = (1\Omega)I_1 + (2\Omega)I_3 \ldots\ldots (ii)$$

Applying it to the loop containing the second battery and the resistor, with $r_2 = 1\Omega$:

$$E_2 = I_2 r_2 + I_3 R \;\Rightarrow\; 7V = (1\Omega)I_2 + (2\Omega)I_3 \ldots\ldots (iii)$$

Solving the three equations gives:

$$I_1 = 2A, \qquad I_2 = 1A, \qquad I_3 = 3A$$

The first battery supplies $2A$, the second supplies $1A$, and the external resistor carries their sum, $3A$.

Example 55

(a) Joule's law of heating and Ohm's law are sometimes confused. Explain what each law describes and how the two differ.
(b) A battery of EMF $6V$ and internal resistance 0.5Ω is short-circuited by a length of wire of negligible resistance. Determine the current that flows and the rate at which heat is produced inside the battery, and comment on the hazard.

Solution

(a) Ohm's law relates current to voltage: it states that the current through an ohmic conductor is directly proportional to the potential difference across it at constant temperature, so that $V = IR$. Joule's law of heating, by contrast, relates power to

current: it states that the rate at which a resistor converts electrical energy into heat is proportional to the square of the current, so that $P = I^2R$. The two therefore answer different questions, Ohm's law telling us how much current flows and Joule's law telling us how much heat that current produces, and while Ohm's law holds only for ohmic conductors, Joule's law holds for any resistor whatever.

(b) With the external resistance negligible, the only resistance limiting the current is the internal resistance of the battery, so the short-circuit current is:

$$I = I_{short} = \frac{E}{r} = \frac{6V}{0.5\Omega} = 12A$$

The rate of heat production inside the battery is the Joule heating in its internal resistance:

$$P = I^2r = (12A)^2(0.5\Omega) = 72W$$

A current of $12A$ flows and $72W$ is produced inside the battery. This is dangerous: the whole of the battery's power is now dumped as heat within the cell itself, so it overheats rapidly and may leak, rupture, or explode, which is why a short circuit must never be allowed.

Example 56

(a) In everyday speech the words "energy" and "power" are used loosely. Distinguish between electrical energy and electrical power as the terms are used in physics, and give the SI unit of each.

(b) A $2.0kW$ electric kettle heats $1.5kg$ of water from $25°C$ to $100°C$. Assuming no heat is lost to the surroundings, determine the time taken. Take the specific heat capacity of water as $4200Jkg^{-1}K^{-1}$.

Solution

(a) Electrical energy is the total work done by an electric current in transferring or converting charge's energy, and its SI unit is the joule. Electrical power is the rate at which that energy is supplied or converted, that is the energy per unit time, and its SI unit is the watt, where one watt equals one joule per second. Energy is therefore the total amount delivered, whereas power is how fast it is delivered.

(b) The heat the water must absorb follows from the specific heat capacity, with a temperature rise of $\Delta\theta = 100°C - 25°C = 75K$:

$$Q = mc\Delta\theta = (1.5kg)(4200Jkg^{-1}K^{-1})(75K) = 472500J$$

Since the kettle supplies energy at the steady rate of $2.0kW$, the time to deliver this heat is:

$$t = \frac{Q}{P} = \frac{472500J}{2000W} = 236s$$

The kettle takes about $236s$, which is approximately $3.9min$, to bring the water to the boil.

Example 57

(a) Explain what is meant by the figure of merit of a galvanometer, and state how it is related to the sensitivity of the instrument.

(b) A galvanometer gives a full-scale deflection of 25 divisions when a current of $2.0mA$ passes through it. Determine its figure of merit. Given that its coil has a resistance of 50Ω, determine the potential difference across it at full-scale deflection.

Solution

(a) **The figure of merit** of a galvanometer *is the current required to produce a deflection of one division on its scale*. It is the reciprocal measure of sensitivity: a galvanometer with a small figure of merit needs only a tiny current for each division and is therefore highly sensitive, whereas a large figure of merit means a large current is needed per division and the instrument is less sensitive.

(b) The figure of merit is the current divided by the number of divisions it deflects:

$$k = \frac{I}{n} = \frac{2.0 \times 10^{-3}A}{25} = 8.0 \times 10^{-5}A \text{ per division}$$

The potential difference across the coil at full-scale deflection follows from Ohm's law:

$$V = IR = (2.0 \times 10^{-3}A)(50\Omega) = 0.1V$$

The galvanometer has a figure of merit of $8.0 \times 10^{-5}A$ per division, and the potential difference across it at full-scale deflection is $0.1V$.

Example 58

(a) A gas is normally an insulator, yet under the right conditions it will carry a current. Explain the conditions under which a gas conducts, and explain what is meant by ionisation.

(b) A spark first jumps across a $2.0cm$ air gap when the potential difference between the electrodes reaches $60kV$. Determine the breakdown field of the air, and hence find the voltage at which a spark would jump across a $3.0cm$ gap between the same electrodes.

Solution

(a) A gas conducts only when it contains free charge carriers, and a neutral gas has none, so carriers must first be created by ionisation. Ionisation is the removal of one or more electrons from a neutral gas molecule, which leaves a positive ion and a free electron, both of which can carry current. It requires energy, supplied either by a strong electric field or by an ionising agent such as ultraviolet light, X-rays, or radiation, and the gas conducts only once the applied voltage reaches the breakdown value at which ionisation begins.

(b) The breakdown field is the breakdown voltage divided by the gap across which it acts:

$$E = \frac{V}{d} = \frac{60 \times 10^{3}V}{2.0 \times 10^{-2}m} = 3.0 \times 10^{6}Vm^{-1}$$

For the same air the breakdown field is unchanged, so the voltage needed across the wider gap follows from $V = Ed$:

$$V = Ed = (3.0 \times 10^{6}Vm^{-1})(3.0 \times 10^{-2}m) = 90kV$$

The breakdown field of the air is $3.0 \times 10^{6}Vm^{-1}$, and the wider 3.0cm gap requires the larger voltage of 90kV to spark.

Example 59

(a) Distinguish between thermionic emission and photoelectric emission.
(b) In a cathode-ray tube an electron starts from rest and is accelerated through a potential difference of 2.0kV. Determine the kinetic energy it gains, in joules and in electron-volts, and its final speed. Take the electronic charge as $1.6 \times 10^{-19}C$ and the electron mass as $9.1 \times 10^{-31}kg$.

Solution

(a) Both processes release electrons from the surface of a metal, but they differ in the energy that frees them. In thermionic emission the metal is heated, and the electrons gain enough thermal energy to escape from its surface. In photoelectric emission the metal is illuminated, and an electron escapes by absorbing a photon of sufficiently high frequency. The one is therefore driven by heat and the other by light.

(b) The kinetic energy gained equals the work done on the electron by the accelerating voltage:

$$E_k = qV = (1.6 \times 10^{-19}C)(2000V) = 3.2 \times 10^{-16}J$$

Expressed in electron-volts, the energy gained in falling through 2000V is simply:

$$E_k = 2000eV = 2.0keV$$

The final speed follows from $E_k = \frac{1}{2}mv^2$, so that:

$$v = \sqrt{\frac{2E_k}{m}} = \sqrt{\frac{2(3.2 \times 10^{-16}J)}{9.1 \times 10^{-31}kg}} = 2.7 \times 10^{7}ms^{-1}$$

The electron gains $3.2 \times 10^{-16}J$, equal to 2.0keV, and reaches a final speed of $2.7 \times 10^{7}ms^{-1}$.

Example 60

(a) Faraday summarised electrolysis in two laws. Explain what each law tells us, and explain what is meant by the electrochemical equivalent of a substance.
(b) A steady current of 2.0A is passed through copper sulphate solution for 30min. Determine the mass of copper deposited on the cathode. Take the electrochemical equivalent of copper as $3.3 \times 10^{-7}kgC^{-1}$.

Solution

(a) Faraday's first law tells us that the mass of substance liberated at an electrode is directly proportional to the quantity of charge passed through the electrolyte. His second law tells us that, for the same quantity of charge, the mass liberated is proportional to the chemical equivalent of the substance, so that different substances are deposited in proportion to their equivalent weights. The electrochemical equivalent of a substance is the mass of it liberated by one coulomb of charge.

(b) The charge passed is the product of the current and the time, with the time first converted to seconds:

$$Q = It = (2.0A)(30 \times 60s) = 3600C$$

By Faraday's first law the mass deposited is the electrochemical equivalent multiplied by the charge:

$$m = zQ = (3.3 \times 10^{-7}kgC^{-1})(3600C) = 1.2 \times 10^{-3}kg$$

The mass of copper deposited on the cathode is $1.2 \times 10^{-3}kg$, that is about 1.2g.

Example 61

(a) Distinguish between line, band, and continuous spectra, and give one source of each.
(b) A diffraction grating is ruled with 5000 lines per centimetre. Determine the spacing of its lines, and hence the highest order in which light of wavelength 600nm can be observed with it.

Solution

(a) A line spectrum consists of separate bright lines on a dark background and is produced by an excited monatomic gas at low pressure, such as a sodium lamp. A band spectrum consists of groups of closely spaced lines, each group crowding

together into a band, and is produced by an excited molecular gas, such as nitrogen. A continuous spectrum is an unbroken range of colour with no gaps and is produced by a hot dense source, such as the glowing filament of a lamp.

(b) The spacing of the lines is one centimetre divided by the number of lines in it:

$$d = \frac{1 \times 10^{-2}\text{m}}{5000} = 2.0 \times 10^{-6}\text{m}$$

The highest order is fixed by the requirement that $\sin\theta$ cannot exceed 1 in the grating equation $m\lambda = d\sin\theta$, so that $m \leq \frac{d}{\lambda}$:

$$\frac{d}{\lambda} = \frac{2.0 \times 10^{-6}\text{m}}{600 \times 10^{-9}\text{m}} = 3.3$$

Since the order must be a whole number not exceeding 3.3, the highest observable order is the third. The grating has a line spacing of $2.0 \times 10^{-6}\text{m}$, and 600nm light can be seen up to the third order.

Example 62

(a) Using the energy changes that take place in each, explain why an incandescent filament lamp is far less efficient than a fluorescent or LED lamp.

(b) A shop replaces ten 60W incandescent bulbs with ten 8W LED bulbs, each lit for 10h each day. Determine the energy saved in a 30-day month, and the money saved if TANESCO charges 290TZS for each kilowatt-hour.

Solution

(a) In an incandescent lamp the current heats a filament until it glows, so the energy is converted first into heat and only the small part of it that falls in the visible range leaves as light; the great majority is wasted as infrared heat. In a fluorescent or LED lamp the electrical energy is converted into light far more directly, with very little spent on heating, so a much larger fraction of the input energy emerges as visible light. The incandescent lamp is therefore far less efficient, because it is essentially a heater that happens to glow.

(b) The energy used by the ten incandescent bulbs in the month is the total power multiplied by the running time:

$$E_{\text{inc}} = (10 \times 60\text{W})(10\text{h})(30) = 180\text{kWh}$$

The energy used by the ten LED bulbs over the same period is:

$$E_{\text{LED}} = (10 \times 8\text{W})(10\text{h})(30) = 24\text{kWh}$$

The energy saved is the difference between the two:

$$E_{\text{saved}} = 180\text{kWh} - 24\text{kWh} = 156\text{kWh}$$

The money saved is this energy multiplied by the tariff:

$$\text{Saving} = 156\text{kWh} \times 290\text{TZS/kWh} = 45240\text{TZS}$$

The change saves 156kWh and 45240TZS each month.

Example 63

(a) Explain how a fuse protects an electrical circuit, and explain why it is connected in the live wire rather than the neutral.

(b) A 13A fuse protects a 240V socket. An electric iron rated 1.0kW, a kettle rated 2.0kW, and a heater rated 1.5kW are all run from that socket at the same time. Determine the total current drawn and state whether the fuse blows.

Solution

(a) A fuse is a short piece of thin wire of low melting point connected in series with the circuit. If the current rises above the safe value, the heating produced in the fuse melts it, the circuit breaks, and the rest of the wiring is protected from overheating and fire. It is placed in the live wire so that, when it blows, the appliance is cut off from the dangerous live potential; were it placed in the neutral, the appliance would remain connected to the live supply even after the fuse had melted, leaving it a shock hazard.

(b) The three appliances together draw a total power that is the sum of their ratings:

$$P = 1.0\text{kW} + 2.0\text{kW} + 1.5\text{kW} = 4.5\text{kW}$$

The total current drawn from the 240V supply follows from $P = VI$:

$$I = \frac{P}{V} = \frac{4500\text{W}}{240\text{V}} = 18.75\text{A}$$

Since 18.75A exceeds the 13A rating of the fuse, the fuse blows and the circuit is broken.

Example 64

(a) Explain why the appliances in a house are connected in parallel across the mains rather than in series.

(b) Two identical bulbs, each of resistance 6Ω, are connected across a 12V supply of negligible internal resistance, first in series and then in parallel. Determine the power dissipated in one bulb in each arrangement, and state which arrangement makes the bulbs brighter.

Solution

(a) Appliances are connected in parallel so that each receives the full mains voltage and can be switched on and off independently of the others. Were they in series, they would share the supply voltage between them, so each would receive

less than its rated voltage and run dimly or weakly, and a break anywhere in the chain, or the switching off of any one appliance, would interrupt the current to them all. Parallel connection avoids both faults.
(b) In series the two bulbs add to a total resistance of 12Ω, so the current is:

$$I = \frac{V}{R_{series}} = \frac{12V}{12\Omega} = 1.0A$$

and the power in one bulb is:

$$P_{series} = I^2R = (1.0A)^2(6\Omega) = 6.0W$$

In parallel each bulb has the full $12V$ across it, so the power in one bulb is:

$$P_{parallel} = \frac{V^2}{R} = \frac{(12V)^2}{6\Omega} = 24W$$

Each bulb dissipates $6.0W$ in series but $24W$ in parallel, so the bulbs are brighter in the parallel arrangement.

Example 65

(a) Explain the difference between an initiated and a self-sustaining gas discharge, and explain the part that chain ionisation plays in the self-sustaining case.
(b) The gas in a discharge tube carries a current of $8.0mA$ when a potential difference of $800V$ is maintained across it. Determine the power dissipated in the tube and the energy it delivers to the gas in $5.0min$.

Solution

(a) An initiated discharge conducts only while an external ionising agent keeps supplying ions, so that the current dies the moment the agent is removed; the applied voltage merely sweeps the ready-made ions to the electrodes. A self-sustaining discharge continues on its own without any external agent, because the voltage is now large enough for the discharge to create its own ions. The mechanism that makes this possible is chain ionisation: each accelerated electron gains enough energy between collisions to ionise the molecule it strikes, the freed electrons are accelerated in turn and ionise further molecules, and this avalanche keeps the supply of carriers going without outside help.

(b) The power dissipated in the tube is the product of the voltage and the current:

$$P = VI = (800V)(8.0 \times 10^{-3}A) = 6.4W$$

The energy delivered is the power multiplied by the time, with the time converted to seconds:

$$E = Pt = (6.4W)(5.0 \times 60s) = 1920J$$

The tube dissipates $6.4W$ and delivers $1920J$, that is $1.92kJ$, to the gas in five minutes.

Sixteen miscellaneous worked examples now lie behind you, and between them they have called on every corner of the chapter, from the drift of an electron to the breakdown of air, and demanded that you tell each one apart by sight. If they have left you wanting more, you are in luck, and if they have left you wanting mercy, you are out of it: the Digging Deeper questions that follow give the scenario and the instruction and nothing else, and this time there is no worked solution waiting at the foot of the page. Pen ready.

DIGGING DEEPER EXERCISE 12

EXERCISE 12A: BINDER QUESTIONS

Question 1

Explain why the drift velocity of electrons in a copper wire is of the order of a millimetre per second, while the signal that switches on a lamp travels almost at the speed of light. Use the queue-and-buckets argument.

Question 2

Distinguish between the EMF and the terminal voltage of a cell, and give the condition under which the two are equal.

Question 3

Ohm's law is often quoted in a single sentence. Express it in its simplest form, identify two materials that obey it and two that do not, and explain why it is described as empirical rather than as a fundamental law of nature.

Question 4

Explain what is meant by the conductance of a conductor, give its SI unit, and suggest why conductance is seldom used in school-level circuit problems.

Question 5

Explain why the resistance of a metal increases with temperature, referring both to the relaxation time and to the microscopic expression.

Question 6

Distinguish between resistance and resistivity, and give the SI unit of each.

Question 7

Explain why nichrome or constantan, rather than copper, is used for heating elements and standard resistors.

Question 8

Kirchhoff gave two laws for electrical networks. Explain what each one says, and name the conservation principle that each embodies.

Question 9

Explain why a Wheatstone bridge gives a more accurate value of a resistance than a direct ohmmeter reading.

Question 10

Explain why the potential gradient along a potentiometer wire must be uniform, and give the conditions on the wire and on the driving current that guarantee it.

Question 11

Explain how a fuse made of a low-melting-point alloy protects a household circuit.

Question 12

Explain what the maximum power transfer theorem says, and why power transmission systems such as TANESCO's national grid deliberately avoid the maximum-power condition.

Question 13

Distinguish between an initiated and a self-sustaining gas discharge, and explain the part that chain ionisation plays in the self-sustaining regime.

Question 14

Explain why a sodium street lamp glows yellow-orange while a neon advertising sign glows red. Refer to the discrete energy levels of the gas atoms.

EXERCISE 12B: REAL QUESTIONS

Question 15

Explain why the metal casing of an electric appliance, such as a refrigerator, is connected to the earth pin of its three-pin plug.

Question 16

TANESCO transmits power between its substations at 132kV but steps it down to 240V before it reaches a house. Explain why such a high voltage is used for long-distance transmission rather than the household voltage.

Question 17

Explain why a torch bulb almost always blows at the instant it is switched on, rather than after it has been glowing steadily for some time.

Question 18

A torch fitted with dry cells is noticeably brighter when the cells are new than after a week of use. Explain this, referring to the EMF, the internal resistance, and the terminal voltage of the cells.

Question 19

Explain why a lamp designed to run on the 240V mains burns out almost at once if it is mistakenly connected to a supply of much higher voltage.

Question 20

Explain why it is dangerous to run a high-power electric heater from a long, thin extension lead that has been left coiled up on its reel.

Question 21

A student measures the resistance of a copper coil with a multimeter and notices that the reading slowly increases the longer the meter is left connected. Explain this observation.

Question 22

A Geiger-Müller counter clicks more rapidly when it is brought closer to a radioactive source. Explain this, referring to the operating region of the ionisation curve.

Question 23

The arc of an electric welder appears bluish-white at its centre but more orange towards its edges. Explain this, referring to the temperature across the arc and the light each part gives off.

Question 24

Explain why the cathode filaments at the two ends of a fluorescent tube are briefly heated before the main discharge begins.

Question 25

A person who touches a faulty switch receives a more severe shock when their hands are wet than when they are dry. Explain why.

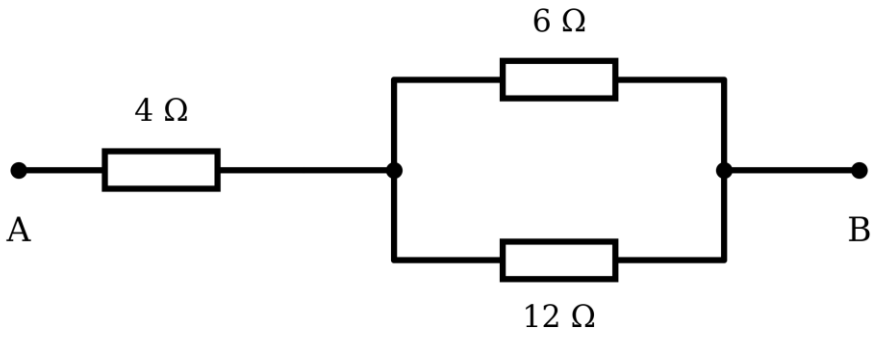
EXERCISE 12C: HOT QUESTIONS

Question 26

A uniform wire of resistance 12Ω is cut into three equal pieces, and the three pieces are then joined side by side in parallel. Determine the resistance of the combination.

Question 27

For the network of resistors in the figure, determine the equivalent resistance between the terminals A and B.



Question 28

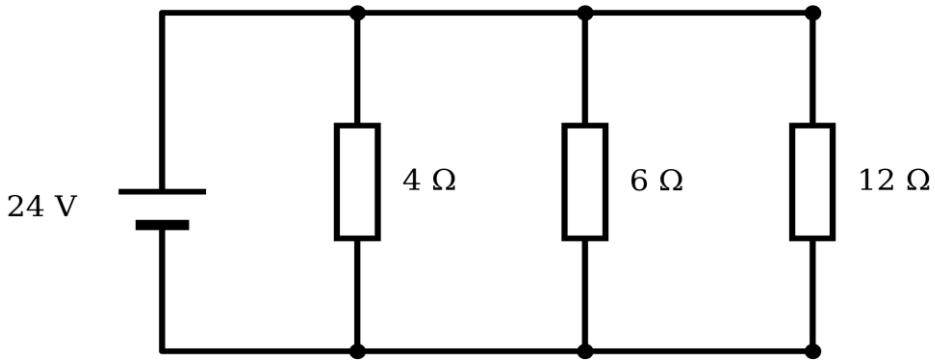
A coil of copper wire has a temperature coefficient of resistance of $4.0 \times 10^{-3}\,^{\circ}\text{C}^{-1}$, referred to its resistance at 0°C. Determine the temperature at which its resistance is 20% greater than its value at 0°C.

Question 29

When a resistor of 4Ω is connected across a cell, the current is 1.0A; when the 4Ω resistor is replaced by one of 9Ω, the current falls to 0.5A. Determine the EMF and the internal resistance of the cell.

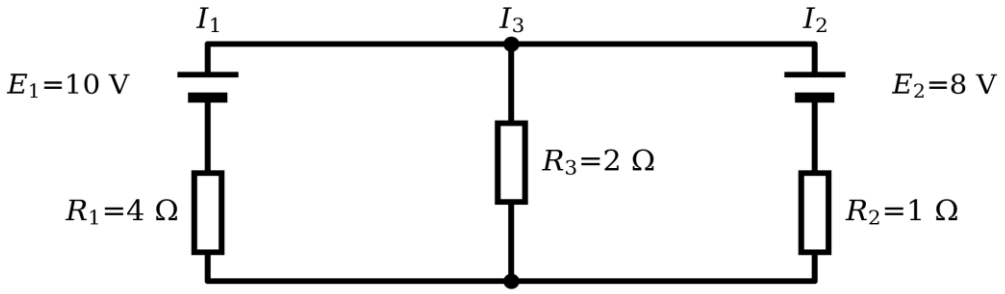
Question 30

Three resistors of 4Ω, 6Ω and 12Ω are connected in parallel across a 24V supply of negligible internal resistance, as shown. Determine the equivalent resistance of the combination and the current in each resistor.



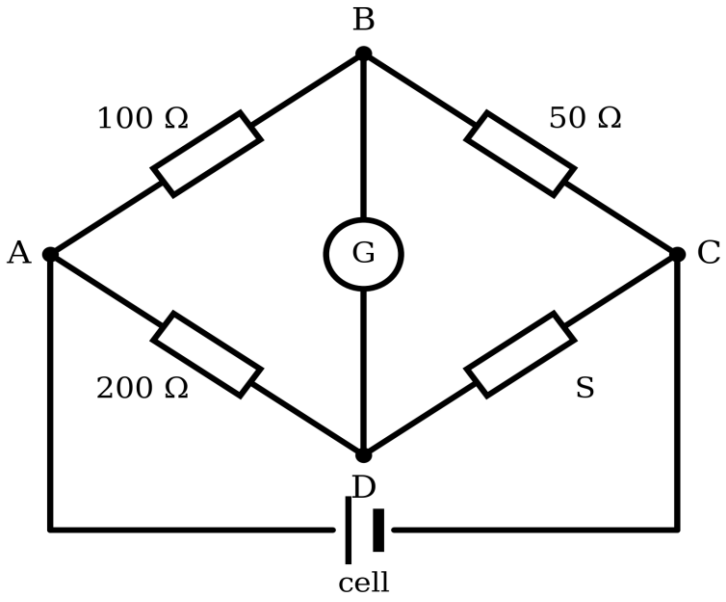
Question 31

In the network of the figure the cells have negligible internal resistance. Using Kirchhoff's laws, determine the current supplied by each cell, the current in the central resistor R_3, and the power dissipated in R_3.



Question 32

The Wheatstone bridge in the figure is balanced, so that no current flows through the galvanometer. Determine the value of the unknown resistance S.

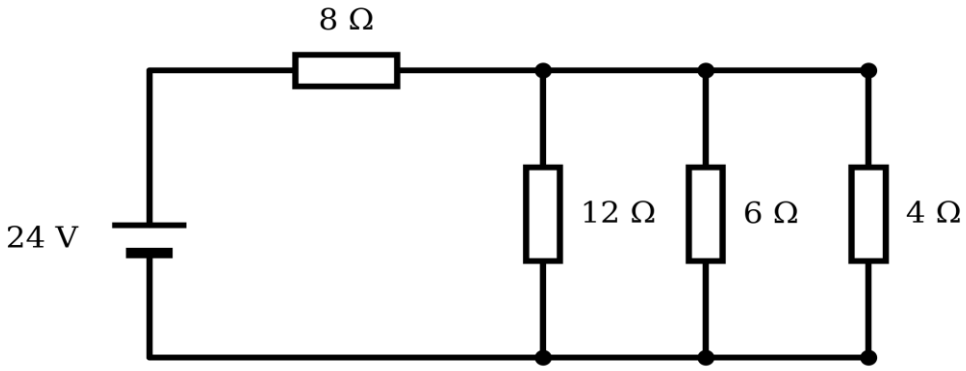


Question 33

Four identical cells, each of EMF 1.5V and internal resistance 0.5Ω, are connected to an external resistor of 2.0Ω, first all in series and then all in parallel. Determine the current through the external resistor in each arrangement.

Question 34

The 24V battery in the figure has negligible internal resistance. Determine the total current drawn from the battery and the current in each of the parallel resistors.



Question 35

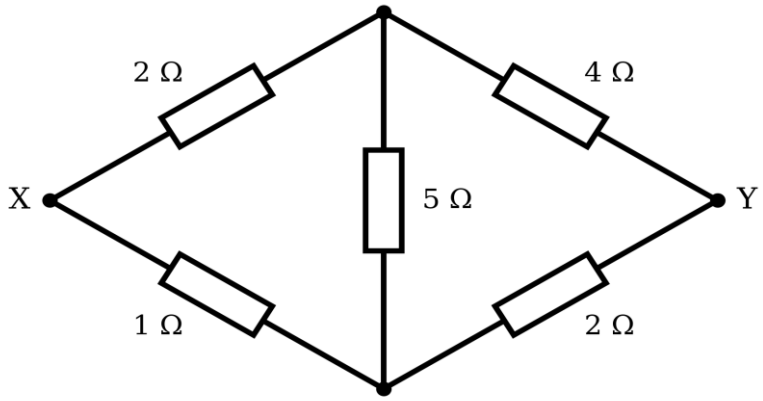
An aluminium wire of cross-sectional area $1.0 \times 10^{-6}\text{m}^2$ carries a steady current of 5.0A. Taking the number of free electrons per unit volume in aluminium as $1.8 \times 10^{29}\text{m}^{-3}$ and the electronic charge as $1.6 \times 10^{-19}\text{C}$, determine the drift velocity of the electrons.

Question 36

Three resistors of 2Ω, 3Ω and 5Ω are joined in series across a supply of 20V with negligible internal resistance. Determine the current in the circuit and the power dissipated in each resistor.

Question 37

For the bridge network in the figure, determine the equivalent resistance between the terminals X and Y.

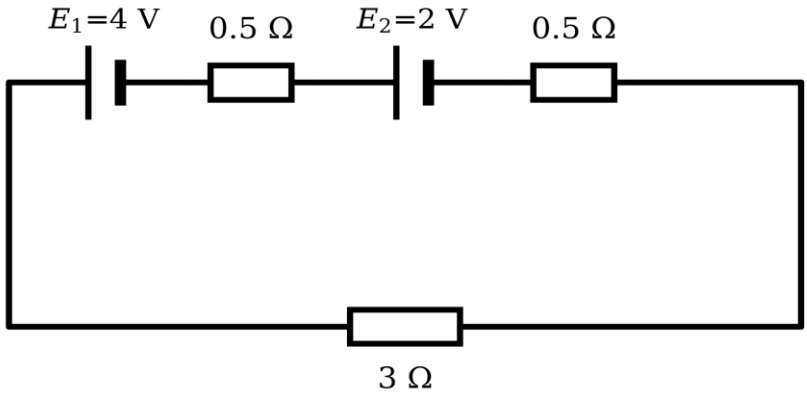


Question 38

Two resistors give a combined resistance of 10Ω when joined in series and 2.4Ω when joined in parallel. Determine the resistance of each.

Question 39

The two cells in the figure each have an internal resistance of 0.5Ω and are joined in series with a 3Ω resistor. Determine the current in the circuit when the cells are connected so that they assist each other, and again when one is reversed so that they oppose each other.

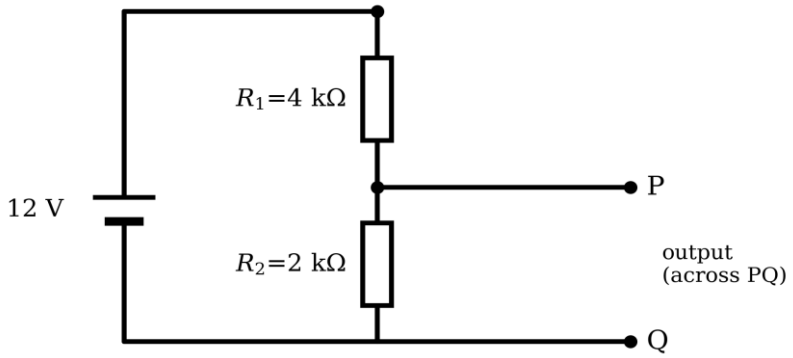


Question 40

A 6Ω resistor and a 3Ω resistor are connected in parallel, and the combination dissipates a total power of 18W. Determine the supply voltage across the combination and the current in each resistor.

Question 41

In the potential divider of the figure a 12V supply is connected across a 4kΩ resistor in series with a 2kΩ resistor, and the output is taken across the 2kΩ resistor at the terminals P and Q. Determine the output voltage when nothing is connected across PQ, and again when a load of 2kΩ is connected across PQ.

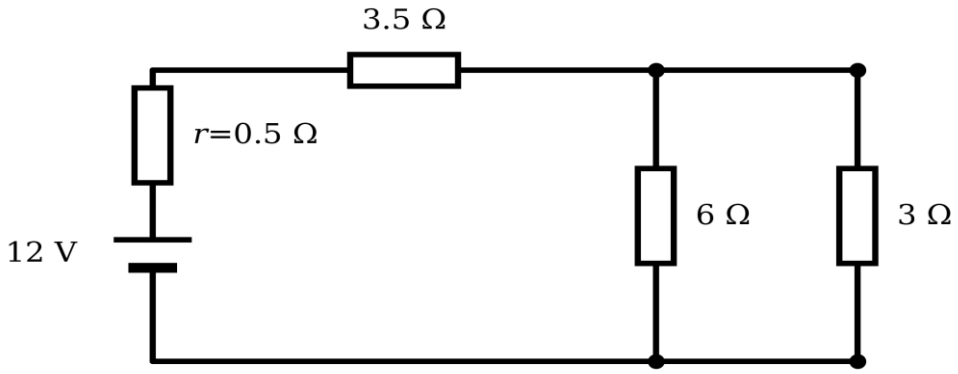


Question 42

A single stroke of lightning transfers a charge of 15C from the cloud to the ground in a time of 0.20ms, through a potential difference of 3.0×10^8V. Determine the average current in the stroke, the energy released, and the average power delivered.

Question 43

The cell in the figure has an EMF of 12V and an internal resistance of 0.5Ω, and it drives current through a 3.5Ω resistor in series with a parallel combination of a 6Ω and a 3Ω resistor. Determine the current drawn from the cell, the terminal voltage of the cell, the power delivered to the external circuit, and the power wasted inside the cell.



Chapter 13

ALTERNATING CURRENT THEORY

INTRODUCTION

On the last afternoon of Chapter 12, with the steady current of batteries and bulbs still warm in everyone's memory, Mr. Akilikubwa made a tactical error. He asked the class an open question. He wanted to know which they thought was the finer kind of electricity: the loyal, one-way current they had just spent a whole chapter studying, or the restless current that comes out of the wall. He really should have known better than to hand Kipanga an opening for a speech.

Kipanga rose to his feet, holding up his pocket torch like a small trophy.

Kipanga: *Sir, there is no contest. In Chapter 12 I met a current with **character.** It chooses a direction and it commits. The electrons in this torch set out from the battery, crawl to the bulb and back, slower than an ant carrying a leaf, the very same way every second of every day. That is a current a man can trust. But the current from the wall? It sets off, loses its nerve, turns around, comes back, and changes its mind again, fifty times a second! That is not a current, sir. That is a **quarrel.** Why would TANESCO build an entire country on a current that cannot decide whether it is coming or going?*

Kipute did not even look up from her notebook.

Kipute: *Because the indecision is the entire point, Kipanga. Your faithful torch current can never be lifted to a high voltage and carried across the country. The quarrelsome one can. That is the only reason the wires above your house are humming while your torch sits quiet in your hand.*

Mr. Akilikubwa: *Kipute has gone straight to the heart of it. In Chapter 12 we told you, and asked you simply to trust us, that TANESCO carries its power across Tanzania at hundreds of thousands of volts, then steps it gently back down to the 240 at your socket. We never showed you the machine that does the stepping. That machine is the **transformer**, and it keeps one stubborn condition: it works only on a current that will not hold still. Offer it the dependable, committed current of Kipanga's torch and it does precisely nothing: no stepping up, no stepping down, only a coil of wire sitting there uselessly warm. A hundred years ago the engineers of the world fought bitterly over this very question, the one Kipanga has just settled so confidently in his own favour. The restless current won. It won for a single reason, the transformer, and that reason is the quiet hero of this chapter.*

Kipanga, unwilling to surrender either his torch or his thesis, tried one last line of defence.

Kipanga: *But sir, in Chapter 11 the capacitor was perfectly content to sit still and hold its charge for a week, long enough, as my thumb has not forgotten, to bite me through the back of a television that was not even plugged in. Surely the components themselves prefer a current that behaves.*

Mr. Akilikubwa (with the faint smile of a man about to spring a trap): *That same capacitor, Kipanga, fed this restless current, will never sit still again. It will fill and empty fifty times a second, and in doing so it will pass a current it could never have passed from your torch. Fully half of what the last two chapters trained you to expect is about to misbehave. I suspect you are going to enjoy the misbehaviour a great deal more than you currently intend to.*

Kipanga sat down, still holding the torch, but eyeing the wall socket with the wary respect of a man who has just been told that the quiet one in the room is the dangerous one.

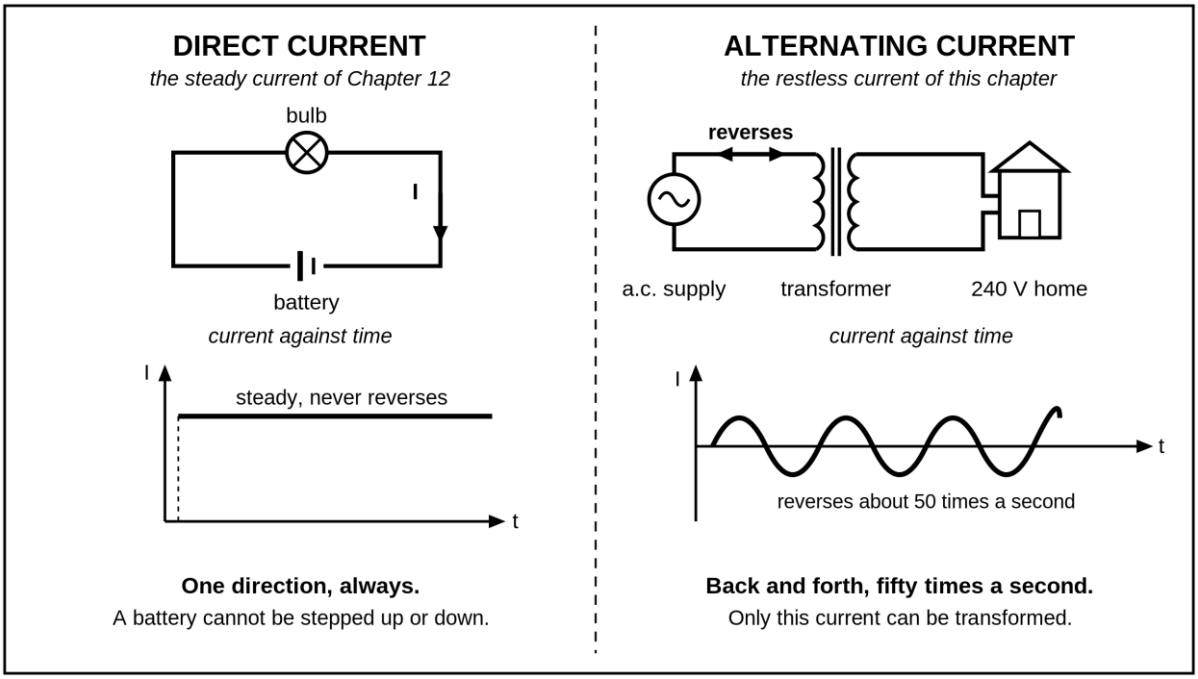


Figure: ***Two currents, side by side.*** *On the left, the steady one-way current of a torch: a battery drives a fixed current through the bulb in one direction only, and the current against time is a flat line that never reverses. On the right, the supply from the wall reverses about fifty times each second, and it feeds a transformer that steps the voltage gently down to the 240V of a home. Only the restless current on the right can be stepped up or down, which is the whole reason the grid is built from it.*

Two chapters stand directly behind this one, and alternating current unsettles them both. In **Chapter 11** a capacitor charged once, through its resistor, in a handful of time constants, and then held its charge in perfect stillness, long enough to sting a careless hand reaching into an unplugged set. In **Chapter 12** the current settled into one patient direction and kept it forever, from the torch in a student's pocket to the immersion heater in a Nyamagana kitchen.

Alternating current keeps neither promise. Give a capacitor an alternating voltage and it never finishes charging. Drive a wire from an alternating source and the patient drift of Chapter 12 never commits to a direction at all; it sloshes back and forth and arrives nowhere. Almost every steady habit those two chapters taught us must now be re-examined for a current that simply will not stand still. And, as Mr. Akilikubwa promised, most of it becomes more interesting for the unrest, not less.

This chapter builds that new physics in the order it is best learned. We begin where the alternating current is born, in a coil turned inside a magnetic field, and show why the EMF it delivers is bound to be a sine wave. We learn to measure it honestly, through the root-mean-square value hiding behind every "240 V" label. We meet the phasor, a spinning arrow that turns the clumsy trigonometry of alternating quantities into pictures the eye can trust. We then ask how the three building blocks of any circuit answer an alternating voltage: the resistor of O-level, the capacitor of Chapter 11, and the inductor we are meeting properly for the first time. We combine them in series and in parallel, and we arrive at resonance, the knife-edge that lets a small radio pick one station out of the crowded air. We build the transformer in full and discover why the grid simply had to be alternating. We settle the accounts of AC power, and find out why a Dar es Salaam workshop can be charged for power it never used. And we close by taming the restless current back into the steady one of Chapter 12, with the humble rectifier hidden inside every phone charger, radio, and laptop in the country.

Welcome to Chapter 13. Kipanga is about to spend an entire chapter discovering that the most useful current in the world is the one that cannot make up its mind. And he, of all people, should have been the first to see the appeal.

ALTERNATING EMF AND CURRENT

The introduction left us with a current that changes its mind fifty times a second. That is a fine thing to say in words, but physics asks two harder questions, and this section answers both. *First, what does that restless current actually look like when we write it down honestly? Second, where does it come from in the first place?* By the end of this section you will be able to draw the wave, name every feature on it, and build its shape from nothing grander than a coil turning in a magnetic field.

Picturing a current that will not sit still

Imagine watching the voltage at a wall socket on a screen that plots its value against time. It refuses to hold at one level the way the torch of Chapter 12 did. Instead it climbs smoothly from zero up to a greatest value in one direction, eases back down through zero, sinks to an equally great value in the opposite direction, and climbs back to zero again. Then it repeats, with the patience of a tide. That smooth, symmetric, endlessly repeating swing is a **sine wave**, and it is the gentlest back-and-forth that nature allows.

This behaviour is what we now call alternating current. An **alternating current (a.c.)** *is an electric current that periodically reverses its direction:* it grows to a maximum one way, falls back to zero, reverses to an equal maximum the other way, and returns, over and over in a regular cycle.

The source that drives it, an **alternating EMF**, *is one whose magnitude and polarity change periodically with time*. When that change follows a sine curve, as it does for the supply from a rotating generator, the EMF and current are said to be **sinusoidal**, and sinusoidal a.c. is the kind this chapter studies. (A steady torch current, by contrast, is **direct current**: one fixed direction, never reversing.)

Two quantities follow this same swing: the alternating EMF of the generator, and the alternating current it drives. Writing **E** for the **instantaneous EMF**, the value at one chosen instant, and $\mathbf{E_0}$ for the **peak value**, the height of the swing and the exact counterpart of the amplitude from Chapter 6, the whole wave is captured by $\boldsymbol{E = E_0 sin\omega t}$. The current keeps the same law, $\boldsymbol{I = I_0 sin\omega t}$, where I_0 is the peak current, and I is the **instantaneous current** (to distinguish it from I, which is commonly used for a steady or constant current, the lowercase letter $\boldsymbol{i}$ is used to denote instantaneous current).

The sine, the peak, and the rhythm of the swing are all we need. The voltage at a household socket swings between $+339$V and -339V, fifty times every second. Why we politely call this 240V rather than 339V is the business of the next section.

Period, frequency and angular frequency

Three numbers describe the rhythm, and they are really three ways of saying one thing.

The **period** T is the time for one complete swing, measured in seconds. The **frequency** f is the number of complete swings in one second, measured in hertz. They are reciprocals of one another, $T = \frac{1}{f}$. Tanzania's mains supply runs at $f = 50\text{Hz}$, so one full swing takes $T = \frac{1}{50\text{Hz}} = 0.02\text{s}$, a fiftieth of a second. A single blink of your eye contains about five complete cycles of the current lighting the room.

The third number, the **angular frequency** ω, is the one students meet last and fear most, quite unnecessarily. The sine function does not count swings; it counts radians, and one complete swing carries the wave through 2π radians. So the angular frequency is simply how many radians of the cycle go by each second, $\omega = 2\pi f$, measured in radians per second. For our mains, $\omega = 2\pi \times 50\text{Hz} = 314\text{rad s}^{-1}$. You can also read the period straight back from it, since $T = \frac{2\pi}{\omega}$.

Keep all three in view: f counts cycles for the engineer, T counts seconds for the clock, and ω counts radians for the sine. They never disagree with one another.

If all of this feels familiar, it should. It is the very mathematics of Chapter 6, where a mass on a spring and Kipute on her swing rose and fell as $x = A\sin\omega t$. Everything carries straight across. The peak value E_0 here plays exactly the part the amplitude A played there; the period T, the frequency f, and the angular frequency $\omega = 2\pi f$ are defined identically and obey the same relations; and the sine is the same sine. Only the oscillating quantity has changed. In Chapter 6 it was a displacement in space, while here it is a voltage and a current in a wire. The travelling waves of Chapter 7 told the same story with $y = A\sin\omega t$. Learn the alternating-current sinusoid once and you have, honestly, already learned it three times.

One thread is worth tying now. In Chapter 6, pushing the swing at exactly its natural rhythm made its amplitude grow dramatically, and we called that resonance. That same idea returns near the end of this chapter, where tuning a circuit to the right rhythm is precisely how a radio picks one station out of the crowded air.

Where the sine comes from: a coil turning in a field

We promised to earn the sine wave, not merely assert it. The machine that produces alternating current is the alternating-current generator, and Mtera's turbines turn day and night to keep ours running. Before writing a single symbol, let us first picture the arrangement shown in the next figure.

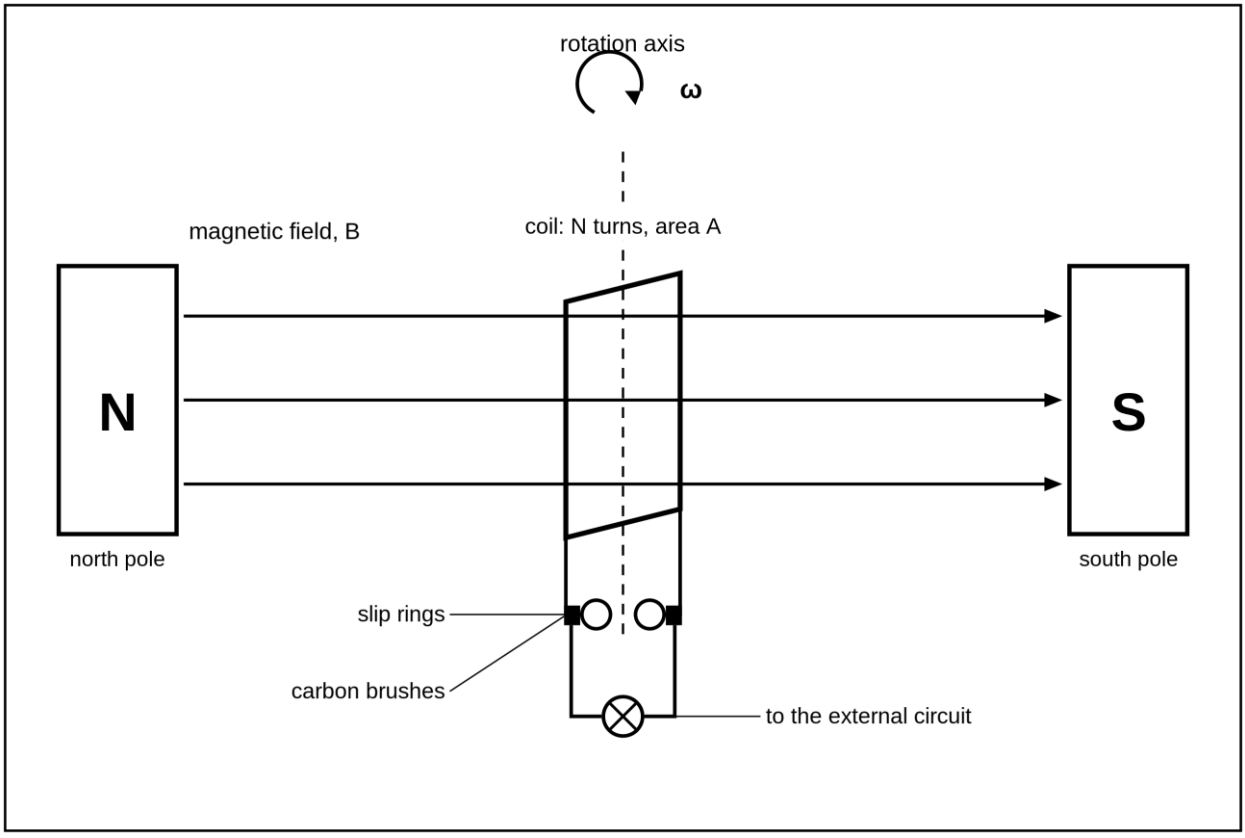


Figure: ***An alternating-current generator****. A coil of N turns and area A is spun at a steady angular velocity ω about an axis lying across a uniform magnetic field B. The turning coil feeds the outside world through two slip rings pressed by carbon brushes.*

A flat coil of N turns, each enclosing an area A, is mounted so that it can spin about an axis lying across a uniform magnetic field B. The turbine spins it at a steady angular velocity ω. As the coil turns, the amount of field threading through it, the **magnetic flux**, rises and falls. *A changing magnetic flux linked with a coil produces an induced EMF.* The goal of the derivation that follows is to convert this physical statement into a quantitative formula.

Start with the flux.

When the face of the coil looks straight along the field, the field pours through it squarely and the flux is greatest. A quarter-turn later the coil's plane lies along the field, no lines thread it, and the flux is zero. In between, the flux follows the cosine of the angle the coil has turned, and because the coil turns at a steady rate that angle is just ωt. So the flux through one turn is $\Phi = BA\cos\omega t$, and for all N turns together the **flux linkage** is:

$$\Lambda = N\Phi = NBA\cos\omega t$$

Now apply Faraday's law.

The induced EMF is the rate at which the flux linkage changes, carrying a minus sign for Lenz's law:

$$E = -\frac{d\Lambda}{dt} = -\frac{d(NBA\cos\omega t)}{dt} = NBA\omega\sin\omega t$$

Set this beside the wave we first drew, $E = E_0\sin\omega t$, and the peak EMF names itself:

$$E_0 = NBA\omega$$

Look at what just happened, drawn out in the following figure.

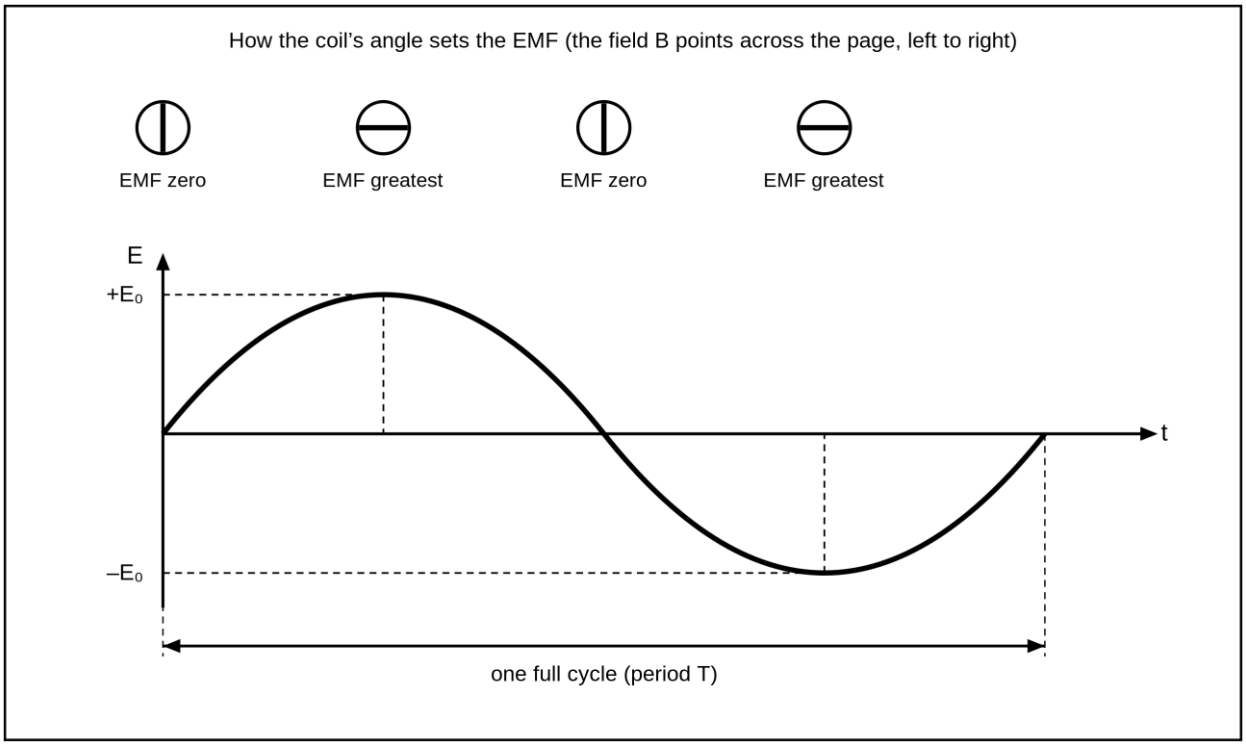


Figure: ***The EMF traces one full sine wave as the coil makes one full turn.*** *The EMF is greatest when the coil sweeps edge-on through the field (flux changing fastest) and zero when the coil faces the field squarely (flux greatest, but momentarily not changing).*

The EMF is greatest at exactly the moments the flux is changing fastest, which is when the coil sweeps edge-on through the field and its wires cut the field lines at full speed. The EMF falls to zero at the moments the flux is greatest, when the coil faces the field squarely and, just for that instant, is not changing at all. This is the part students reliably get backwards, so hold tight to the picture: *greatest EMF where the coil moves **along** the field, zero EMF where it **faces** the field*. The sine wave is not a decoration we chose. It is the unavoidable fingerprint of steady rotation, because the flux varies as a cosine and the rate of change of a cosine is a sine.

The result $E_0 = NBA\omega$ hands you four honest ways to raise the peak voltage:

- wind on more turns N,
- use a stronger magnet B,
- enlarge the coil's area A, or
- spin it faster ω.

Notice that ω quietly does two jobs at once: it sets both the size of the peak and the frequency of the swing. Keep that in mind, because it is the whole engine of the Mtera example coming up. (At Mtera each generator actually carries three coils set a third of a turn apart, which is why the national supply is three-phase; but one coil is all we need for the physics, and all the examination will ask of you.)

One promise for later: this very idea, a changing flux inducing an EMF in a coil, returns near the end of the chapter as the secret of the transformer. Keep it warm.

The theory is on its feet; time to make it work for a living. Five examples follow. Kipanga lobbied hard for two, on the grounds that he had already met enough Greek letters for one afternoon. He was outvoted by the arithmetic.

BINDER Example 1

Tanzania's mains supply has a frequency of 50Hz and a peak voltage of 339V. (a) Find its period. (b) Find its angular frequency. (c) Write the equation for the instantaneous voltage, and use it to find the voltage 0.002s after the voltage last passed through zero on its way up.

Solution

(a) The period is the reciprocal of the frequency:

$$\text{T} = \frac{1}{\text{f}} = \frac{1}{50\text{Hz}} = 0.02\text{s}$$

(b) The angular frequency follows from the frequency:

$$\omega = 2\pi\text{f} = 2\pi \times 50\text{Hz} = 314\text{rad s}^{-1}$$

(c) The instantaneous voltage obeys $E = E_0 \sin\omega t$, so:

$$E = 339\sin 314t$$

Putting $t = 0.002\text{s}$ (so that $\omega t = 314 \times 0.002 = 0.628\text{rad}$):

$$E = 339\sin(0.628)\text{V} = 339 \times 0.588\text{V} = 199\text{V}$$

Making Sense of the Answer: *The period came out as a fiftieth of a second, which is exactly what* $50Hz$ *promises, so the first answer checks itself. The instant we asked about is only a tenth of the way into the cycle, so the voltage has not yet climbed to its crest;* $199V$ *sits comfortably below the* $339V$ *peak, just where it ought to.*

Think Like a Physicist: *Whenever a frequency appears, expect* $2\pi f$ *to walk into every equation behind it, because the sine is fed radians, not cycles. Convert to ω once at the very start and the rest of the problem flows downhill.*

BINDER Example 2

An alternating current in a lamp has a maximum value of 4.0A and a frequency of 50Hz. (a) Write the equation for its instantaneous value. (b) Find the current 0.0025s after it last passed through zero on its way up. (c) Find the time it first takes to reach its peak.

Solution

(a) The instantaneous current obeys $I = I_0 \sin\omega t$, and the angular frequency is :

$$\omega = 2\pi f = 2\pi \times 50\text{Hz} = 314\text{rad s}^{-1}$$

With $I_0 = I_{max} = 4.0\text{A}$:

$$I = 4.0\sin 314t$$

(b) Putting $t = 0.0025\text{s}$, the angle reached is $\omega t = 314 \times 0.0025 = 0.785\text{rad}$, hence:

$$I = 4.0\sin(0.785)\text{A} = 4.0 \times 0.707\text{A} = 2.83\text{A}$$

(c) The current first reaches its peak when $\sin\omega t = 1$, that is when $\omega t = \frac{\pi}{2}$. Solving for the time gives

$$t = \frac{\pi}{2\omega} = \frac{\pi}{2 \times 314\text{rad s}^{-1}} = 0.005\text{s}$$

Making Sense of the Answer: *At one eighth of a cycle the current has climbed to* 0.707 *of its peak, about* $2.83A$*, sitting comfortably between zero and the* $4.0A$ *crest. The first peak then arrives a quarter of a period after a zero, at* $0.005s$*, exactly where the sine first reaches one.*

Think Like a Physicist: *The current law* $I = I_0 sin\omega t$ *is the exact twin of the voltage law* $E = E_0 sin\omega t$*: the same ω, the same sine, only the peak symbol changes from* E_0 *to* I_0*. Read a time-to-peak straight from* $\omega t = \frac{\pi}{2}$*, never by guessing.*

BINDER Example 3

A rectangular coil of 200 turns, each of area 0.020m^2, spins at an angular velocity of 314rad s^{-1} in a uniform magnetic field of 0.15T. Calculate the peak EMF it generates.

Solution

The peak EMF is already the subject of $E_0 = \text{NBA}\omega$, so we substitute directly, keeping every unit in place:

$$E_0 = \text{NBA}\omega = 200 \times 0.15\text{T} \times 0.020\text{m}^2 \times 314\text{rad s}^{-1} = 188.4\text{V}$$

Making Sense of the Answer: *A peak of roughly* $188V$ *is the right size for a modest generator, a little below the household crest we have been quoting, which is reassuring rather than surprising.*

Think Like a Physicist: *The formula* $E_0 = NBA\omega$ *has four handles and no hidden traps. If a question changes just one of* N*,* B*,* A *or* ω *and asks what happens to the peak, you rarely need to compute at all; you reason in proportions.*

REAL Example 4

At the Mtera hydroelectric station the engineers can open the gates and spin the generator coils faster. Kipanga, freshly armed with $E_0 = \text{NBA}\omega$, announces that spinning faster simply makes every light in the country brighter, and proposes that TANESCO should therefore spin the turbines as fast as the river allows. Explain (a) what really happens to the peak EMF when the coil spins faster, (b) what happens at the same time to the frequency, and (c) why TANESCO cannot take his advice.

Solution

(a) Kipanga is right about the voltage. In $E_0 = \text{NBA}\omega$, the turns N, the field B and the area A are all fixed once the machine is built, so the peak EMF is directly proportional to the spin rate ω. Spin the coil faster and the crest of the wave rises in exact step. So far, so bright.

(b) But ω is not free to change on its own. The frequency rides on the very same ω, through $f = \frac{\omega}{2\pi}$. Speeding the coil up does not only lift the voltage; in the same breath it raises the frequency of the whole supply. Faster spin means a higher pitch to the national hum.

(c) And there is the trap. Every mains clock, every motor, and every timing circuit in the country is built to expect exactly 50Hz. Wander away from 50Hz and clocks run fast, motors overheat, and equipment misbehaves from Mtera to Mwanza. So the grid is held at 50Hz with real care: the turbines are governed to a fixed ω, and when the country needs more power the engineers bring more generators online or push more current, never more frequency. Kipanga's lights would indeed glow brighter, for the few minutes before half the appliances in Tanzania filed a complaint.

Making Sense of the Answer: *The mistake was an honest one, because ω really does raise the voltage. The catch is that ω cannot be changed in private; it drags the frequency along with it, and the frequency is the one number the grid is least allowed to touch.*

Think Like a Physicist: *When one symbol plays two roles, changing it has two consequences, not one. Here ω sets both the peak, through $E_0 = NBA\omega$, and the rhythm, through $f = \frac{\omega}{2\pi}$. A careful physicist always asks what else a quantity controls before reaching for its dial.*

HOT Example 5

A coil of 120 turns, each of area $2.5 \times 10^{-2}\text{m}^2$, rotates at 1500 revolutions per minute in a uniform field of 0.20T. (a) Find its angular velocity. (b) Find its period. (c) Find the peak EMF. (d) Write the EMF as a function of time. (e) Find the EMF at $\text{t} = 0.01\text{s}$ and comment on the result.

Solution

(a) The coil completes 1500 cycles in every 60s, so the number of cycles each second is:

$$\text{f} = \frac{1500}{60\text{s}} = 25\text{Hz}$$

$$\omega = 2\pi\text{f} = 2\pi \times 25\text{Hz} = 157\text{rad s}^{-1}$$

The angular velocity is therefore 157rad s^{-1}.

(b) The period is the reciprocal of frequency:

$$\text{T} = \frac{1}{\text{f}} = \frac{1}{25\text{Hz}} = 0.04\text{s}$$

(c) Now the peak EMF:

$$\text{E}_0 = \text{NBA}\omega = 120 \times 0.20\text{T} \times 2.5 \times 10^{-2}\text{m}^2 \times 157\text{rad s}^{-1} = 94.2\text{V}$$

(d) The instantaneous EMF then reads

$$\text{E} = \text{E}_0\sin\omega\text{t} = 94.2\sin 157\text{t}$$

(e) At $\text{t} = 0.01\text{s}$, the angle is $\omega\text{t} = 157 \times 0.01 = 1.57\text{rad}$, which is a right angle, so:

$$\text{E} = 94.2\sin(1.57)\text{V} = 94.2\text{V}$$

Comment: the EMF has reached its full peak value, and this is no accident. The time $\text{t} = 0.01\text{s}$ is exactly one quarter of the period $\text{T} = 0.04\text{s}$, and a quarter of a cycle after a zero is precisely where the sine first equals one. A quarter period after passing through zero, an alternating quantity always sits at its crest.

Making Sense of the Answer: *The answer lands squarely on the peak, and the figure drawn earlier shows why: the quarter-period mark is the very top of the sine. Had we chosen a time that was not a neat fraction of the period, the EMF would have fallen somewhere between zero and 94.2V, and only the formula, not the eye, could have told us where.*

Think Like a Physicist: *Tame the units before you tame the physics. Revolutions per minute become radians per second by multiplying by 2π and dividing by 60; once ω is in proper units, every other quantity drops out in a single line.*

Five examples, and not a single bite. Kipanga has stopped flinching at the sight of ω, which we will record as a modest triumph for education and a quiet disappointment for anyone who came hoping for drama.

MEAN AND RMS VALUES

We can now draw the alternating wave and name its peak. But here is an awkward question. Clamp an a.c. ammeter onto the wire feeding Kipanga's lamp and it reads neither the peak nor zero, but some third number that nobody has yet explained. The supply is stamped 240V, and yet we have just watched the voltage shoot up to 339V and down to -339V. *So which number, honestly, is the supply?* This section settles the matter, and the answer turns out to be one of the most useful ideas in all of a.c. theory.

Why the ordinary average is useless

The first instinct is to take the **mean value**, the way we average a set of test scores: add everything up and divide by how much there was. Try it on the current. Over one complete cycle the current is positive for exactly half the time and negative, by an equal amount, for the other half. Every positive instant is matched by a negative one, the two cancel perfectly, and the ***mean value*** *of an alternating current over a complete cycle is exactly zero.*

That is honest arithmetic, and it is also useless. A current whose average is zero still boils Kipanga's water and still turns the meter that bills his father. Plainly, zero is not what the ammeter is reading. The ordinary average has thrown away the very

thing we care about, because it lets the negative half-cycle erase the positive one. We need a cleverer kind of average, one the minus signs cannot cancel.

The value that matters: the root-mean-square

Here is the escape. The reason an alternating current still heats a wire is that *heating does not care about direction.* A current flowing one way heats a resistor exactly as much as the same current flowing the other way, because the power turned into heat is $P = i^2R$, and squaring a negative number makes it positive. Heat is blind to the sign of the current, and that single observation hands us the right kind of average.

Follow the power. At any instant the power delivered to a resistance R is $P = i^2R = I_0{}^2 R\sin^2\omega t$. The graphs below make the whole idea visible before we touch the algebra.

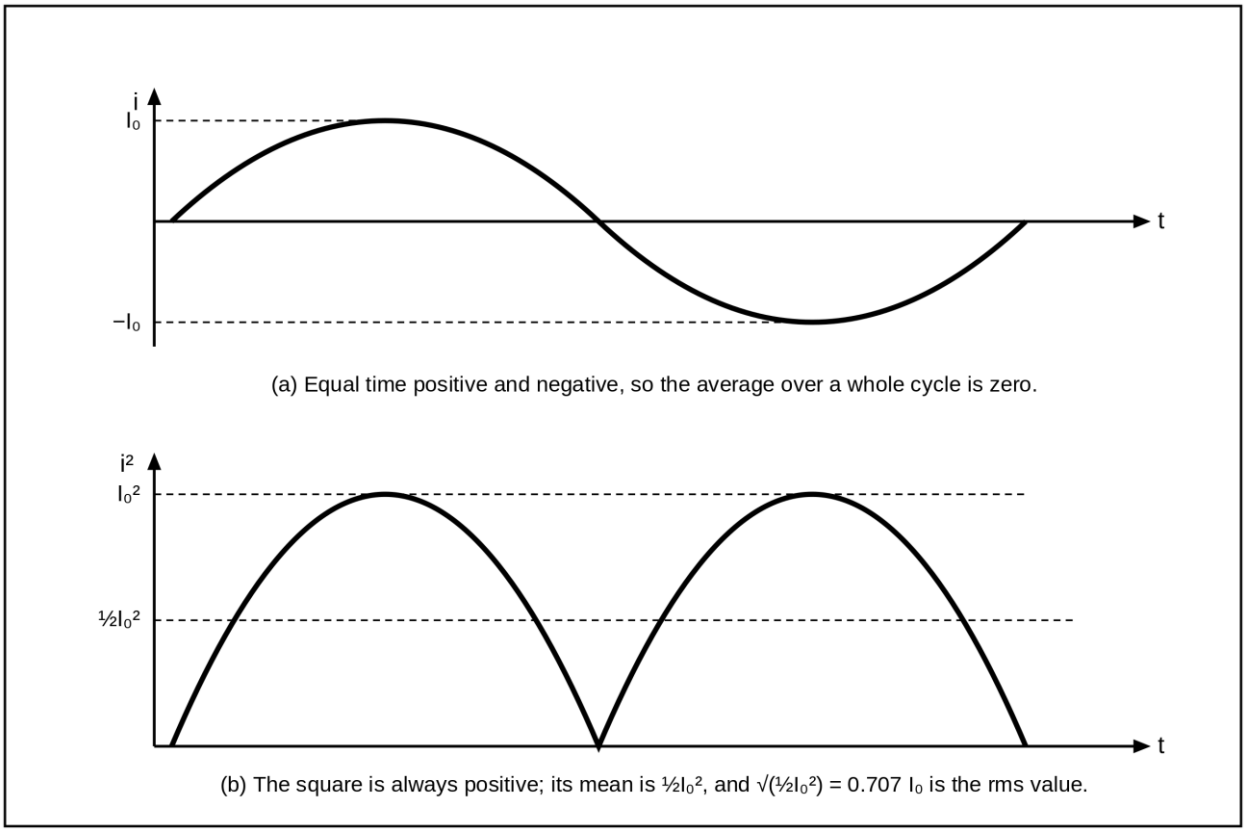


Figure: ***The two averages of an alternating current.*** *The plain average over a full cycle is zero, but the average of the squared current is one half of the peak squared, and the square root of that mean is the root-mean-square value, the steady current that would heat a resistor at the same rate.*

To get the average power we need the average of $\sin^2\omega t$ over a cycle. Here is a clean way to see it: $\sin^2\omega t$ and $\cos^2\omega t$ must average to the same value, since one is only the other shifted along, and at every instant they add to one because $\sin^2\omega t + \cos^2\omega t = 1$. Two equal things that add to one must each be one half. So the average of $\sin^2\omega t$ is ½, and the average power is:

$$P_{av} = \tfrac{1}{2}I_0{}^2 R$$

Now the key question: *what steady, direct current would heat the same resistor at this same rate?* Call it I_{rms}. A steady current delivers power $I_{rms}{}^2 R$, so we set the two equal,

$$I_{rms}{}^2 R = \tfrac{1}{2}I_0{}^2 R$$

The resistance cancels, and the square root gives the headline result,

$$\mathbf{I_{rms} = \frac{I_0}{\sqrt{2}} = 0.707 I_0},$$

and in exactly the same way for the voltage:

$$\mathbf{E_{rms} = \frac{E_0}{\sqrt{2}} = 0.707 E_0}$$

So generally:

$$\mathbf{rms\ value = \frac{peak\ value}{\sqrt{2}}}$$

This number wears a name that describes precisely how it was built. The **root-mean-square value (rms)** of an alternating quantity is found in three steps, in order: **square** the instantaneous value, take the **mean** of the square over a cycle, then take the **square root**. Root, of mean, of square, read backwards. And it carries a physical meaning worth remembering: **The rms value of an a.c.** *is the steady d.c. value that would deliver the same average power, the same heating, to a given resistor.*

If that three-step recipe feels familiar, it should. We used exactly it in Chapter 4, the kinetic theory of gases, where the root-mean-square speed of the molecules stood in for their wildly differing individual speeds. Square, average, take the root. There it tamed a crowd of molecules; here it tames a current that will not sit still. The procedure is identical; only the quantity being averaged has changed.

The mean over half a cycle

There is a second average that occasionally earns its keep, so we name it and move on. If we average the current over just one half-cycle, the positive hump alone, the negatives never get their turn to cancel and the result is not zero. The average height of a single sine arch works out to $\frac{2}{\pi}$ of the peak, so the **mean value over a half cycle** is:

$$\mathbf{I_m = \frac{I_0}{\frac{2}{\pi}} = \frac{2I_0}{\pi} = 0.637I_0,}$$

and likewise:

$$\mathbf{E_m = 0.637E_0}$$

This is a marginal quantity for now, because almost every a.c. calculation reaches for the rms value instead. Its one real moment of usefulness comes later in this chapter, when we pass a.c. through a diode to make d.c.: the steady output of such a rectifier is governed by exactly this half-cycle mean. Until then, keep it in your back pocket.

Form factor

One small ratio (called form factor) ties the two useful averages together. The **form factor** of a waveform *is its rms value divided by its mean value over a half cycle,*

$$\text{form factor} = \frac{I_{rms}}{I_m} = \frac{0.707I_0}{0.637I_0} = 1.11$$

For a pure sine wave this is always 1.11, whatever the peak. It is the number an rms-reading meter must apply internally if what it actually measures is the half-cycle mean, a point we return to when we open up an a.c. meter later in the chapter.

Every rating you read is an rms value

Now let us clear up the confusion that trips nearly every student. *Every a.c. voltage quoted **anywhere**, the* $240V$ *at a Tanzanian socket, the* $11kV$ *and* $400kV$ *on TANESCO's transmission lines, is an **rms value**, never the peak.* Now from:

$$\text{rms value} = \frac{\text{peak value}}{\sqrt{2}};\ \text{peak value} = \sqrt{2} \times \text{rms value}$$

The peak is larger by a factor of $\sqrt{2}$. A 240V socket therefore rises to a peak of $E_0 = \sqrt{2} \times 240V = 339V$, which is why 339V kept appearing in the last section. The distinction is not pedantry. The insulation around a wire and the gap inside a switch must survive the full peak of 339V, so engineers size them on the peak; but the heat in the kettle, the reading on the meter, and the figure on the TANESCO bill all follow the rms value, so those use $240V$. *Peak for safety, rms for power.* Keep the two apart and a great deal of a.c. falls into place.

Three averages, one ammeter reading, and a kettle that refuses to care about signs. Four worked examples turn all of it into numbers. Kipanga asked that none of them contain a square root; we have politely ignored him.

BINDER Example 6

(a) An a.c. supply has a peak voltage of $E_0 = 325V$; find its rms value. (b) A Tanzanian socket has an rms voltage of $E_{rms} = 240V$; find its peak voltage.

Solution

(a) The rms value is the peak divided by $\sqrt{2}$:

$$E_{rms} = \frac{E_0}{\sqrt{2}} = \frac{325V}{\sqrt{2}} = 230V$$

(b) Going the other way, the peak is $\sqrt{2}$ times the rms:

$$E_0 = \sqrt{2} \times E_{rms} = \sqrt{2} \times 240V = 339V$$

Making Sense of the Answer: *The rms is always about* 0.707 *of the peak, and the peak always about* 1.41 *times the rms. Our Tanzanian socket duly returns its familiar* $339V$ *peak.*

Think Like a Physicist: *There are only two conversions here and they are reciprocals: divide by* $\sqrt{2}$ *to go from peak to rms, multiply by* $\sqrt{2}$ *to come back. Decide which way you are travelling before you reach for the calculator.*

REAL Example 7

Kipanga's mother owns two identical 240V, 2.0kW kettles. One is plugged into the mains. The other, during a power cut, runs from his uncle's solar battery bank, which holds a steady 240V. Kipanga bets that the mains kettle, whose voltage is "really 339V at the peak", must boil faster than the battery kettle stuck at a flat 240V. Explain who loses the bet, and why.

Solution

Kipanga loses, and the reason is the whole point of the rms value. The mains voltage does peak at 339V, but it also passes through zero, dwells at every value in between, and spends half of each cycle negative. The single number that captures its heating effect is its rms value, and that rms value is exactly 240V, defined precisely so that the a.c. heats a resistor at the same average rate a steady 240V would. The battery sits at a flat 240V, which is its own rms value, because a steady value is its own root-mean-square.

Both kettles therefore deliver the same average power $P_{av} = \frac{E_{rms}{}^2}{R}$, draw the same rms current, and boil the same litre of water in the same time. The 339V peak is real, but it does its extra work for only an instant and is paid back by the moments near zero. Averaged over a cycle, the two kettles are twins, and Kipanga owes his mother the wager.

Making Sense of the Answer: *The rms value was defined to make exactly this true. If the mains boiled faster than a battery of equal rms voltage, we would have defined the rms value wrongly.*

Think Like a Physicist: *Whenever a problem compares a.c. with d.c. heating, convert the a.c. to its rms value and then treat it as though it were a steady d.c. The rms value is the bridge between the two worlds.*

HOT Example 8

A 2.0kW heating element is connected to the 240V, 50Hz mains, where 240V is the rms voltage. Find (a) the peak voltage, (b) the resistance of the element, (c) the rms current, (d) the peak current, and (e) the average power, confirming it matches the rating.

Solution

(a) The peak is $\sqrt{2}$ times the rms:

$$E_0 = \sqrt{2} \times 240\text{V} = 339\text{V}$$

(b) The element's resistance follows from the average power $P_{av} = \frac{E_{rms}{}^2}{R}$, rearranged for R:

$$R = \frac{E_{rms}{}^2}{P_{av}} = \frac{(240\text{V})^2}{2000\text{W}} = 28.8\Omega$$

(c) The rms current is the rms voltage over the resistance:

$$I_{rms} = \frac{E_{rms}}{R} = \frac{240\text{V}}{28.8\Omega} = 8.33\text{A}$$

(d) The peak current is $\sqrt{2}$ times the rms current:

$$I_0 = \sqrt{2} \times 8.33\text{A} = 11.8\text{A}$$

(e) The average power is the product of the rms voltage and rms current:

$$P_{av} = E_{rms} \times I_{rms} = 240\text{V} \times 8.33\text{A} = 2000\text{W} = 2.0\text{kW},$$

which matches the element's rating, as it must.

Making Sense of the Answer: *The rating "$2.0kW$ at $240V$" is an rms statement, and everything, the resistance, the currents, and the power, flowed from it. The peak current of $11.8A$ is just $\sqrt{2}$ times the rms and plays no part in the power.*

Think Like a Physicist: *Appliance ratings are rms. Begin every appliance problem by reading the plate as rms values, then lean on $P_{av} = E_{rms}I_{rms} = \frac{E_{rms}{}^2}{R}$.*

HOT Example 9

An alternating current is described by $i = 6.0\sin(314t)\text{A}$, with t in seconds. Find: (a) its peak value, (b) its frequency, (c) its rms value, (d) its mean value over a half cycle, (e) its form factor, and (f) the average power it dissipates in a 12Ω resistor.

Solution

(a) Comparing with $i = I_0\sin\omega t$, the peak current is $I_0 = 6.0\text{A}$.

(b) The angular frequency is $\omega = 314\text{rad s}^{-1}$, and since $\omega = 2\pi f$,

$$f = \frac{\omega}{2\pi} = \frac{314\text{rad s}^{-1}}{2\pi} = 50\text{Hz}$$

(c) The rms value is

$$I_{rms} = \frac{I_0}{\sqrt{2}} = \frac{6.0\text{A}}{\sqrt{2}} = 4.24\text{A}$$

(d) The mean over a half cycle is

$$I_m = \frac{2I_0}{\pi} = \frac{2 \times 6.0A}{\pi} = 3.82A$$

(e) The form factor is the ratio of the two,

$$\text{form factor} = \frac{I_{rms}}{I_m} = \frac{4.24A}{3.82A} = 1.11$$

(f) The average power uses the rms current, never the peak:

$$P_{av} = {I_{rms}}^2 R = (4.24A)^2 \times 12\Omega = 216W$$

Making Sense of the Answer: *The form factor came out* 1.11, *exactly as it must for any pure sine, which is a free check on the rms and half-cycle figures. And the power leans on the rms current squared, not the peak.*

Think Like a Physicist: *Read I_0 and ω straight off the equation $i = I_0 sin\omega t$; the frequency, the rms, the mean, the form factor, and the power are each a one-line consequence of those two numbers.*

Four examples, three averages, and one ammeter reading finally explained. Kipanga has made his peace with the square root, mostly because it kept turning the alarming 339V back into the friendly 240V he had trusted all along.

PHASORS: A ROTATING-VECTOR METHOD

So far each of our alternating quantities has lived alone. But the moment we wire two components together we shall have to add their voltages, and here a nasty surprise waits. The voltage across one component may be in step with the current while the voltage across another runs a quarter-cycle ahead, and adding two sine waves that are out of step, by trigonometry, is genuinely unpleasant. Physicists almost never do it. Instead they reach for a picture so simple a child could draw it and so powerful that it turns the whole of a.c. circuit theory into geometry. The picture is the **phasor**, and this short section is given over entirely to it. Master it now, on easy ground, and every circuit in the rest of the chapter becomes a matter of drawing arrows.

A sine wave is the shadow of a spinning arrow

As always, let us begin with an imagination. Imagine an arrow of fixed length, pinned at one end and spinning steadily anticlockwise, like the hand of a clock running backwards at a constant rate. Fix your eye on the height of its tip above the centre line. As the arrow turns, that height climbs from zero up to the arrow's full length, falls back to zero, drops to the full length below the line, and returns. Plot the height against time and you trace a perfect sine wave. A sine wave, in other words, is nothing more exotic than the moving shadow of a spinning arrow, as the following figure shows.

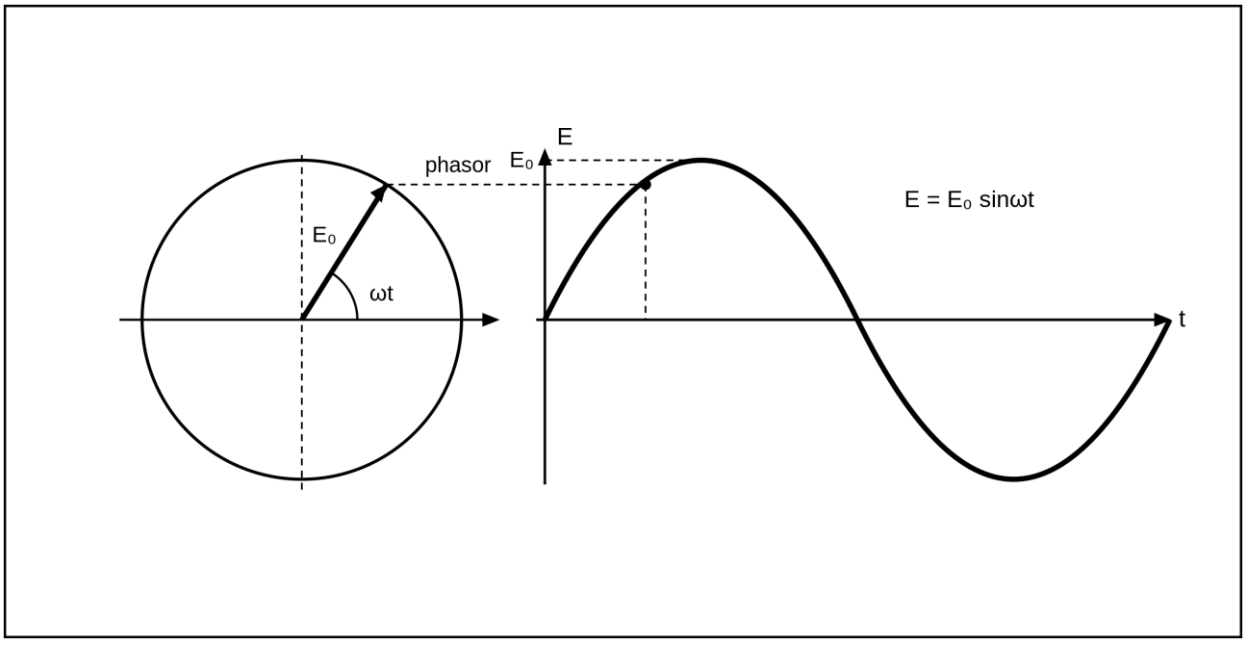


Figure: ***A sine wave is the moving shadow of a spinning arrow.*** *As the phasor of length E_0 turns anticlockwise through the angle ωt, the height of its tip traces out the curve $E = E_0 \sin\omega t$. This is the reference circle of simple harmonic motion, now standing for a voltage.*

You have met this picture before. In Chapter 6 we built simple harmonic motion as the shadow of uniform circular motion: a point ran steadily around a circle and its projection onto a diameter traced out $x = A\sin\omega t$. The **reference circle**, we called it. A **phasor** is that very same rotating radius, borrowed now to stand for a voltage or a current in place of a displacement. The geometry of Chapter 6 has quietly become the bookkeeping of this Chapter.

What a phasor is

Let us pin the idea down. A **phasor** *is a rotating vector that represents a sinusoidal quantity*, and it carries three pieces of information at once:

- its **length** equals the peak value (the amplitude) of the quantity;
- it **rotates anticlockwise** at the angular frequency ω of the quantity; and
- its **projection** onto the vertical at any instant gives the instantaneous value.

A phasor of length E_0, turned through an angle ωt from the horizontal, therefore has a vertical projection $E_0 \sin\omega t$, which is exactly the instantaneous EMF we started from.

We fix a **convention** so that every diagram agrees: *phasors are drawn rotating anticlockwise, with angles measured anticlockwise from a horizontal reference line pointing to the right.* Now comes the liberating fact. *Because every phasor in a given circuit spins at the same frequency ω, the angles between them never change.* We may therefore freeze the rotation at a single instant and draw a still photograph, a **phasor diagram**, in which only two things matter: the length of each arrow, which is its peak value, and the angle between the arrows, which is their phase relationship. The spinning is understood; we never need to draw it twice.

The fixed angle between two phasors is their **phase difference**. *If one phasor sits further round anticlockwise than another, its quantity reaches each stage of the cycle earlier,* and we say it **leads**; *the one behind* **lags**. A phasor a quarter-turn (that is 90°, or $\frac{\pi}{2}$ radians) ahead of another describes a quantity that peaks a quarter-cycle sooner. These two words, leads and lags, will be the entire vocabulary of the circuits to come, where we shall find the voltage across an inductor leading the current and the voltage across a capacitor lagging behind it.

Why the trouble is worth it: adding sinusoids becomes adding arrows

Now the payoff that justifies the whole construction. Two sinusoids of the same frequency, however far out of step, always add to give a third sinusoid of the same frequency. Proving that by trigonometry is a chore; finding it with phasors is a single drawing. Because each sinusoid is the vertical shadow of its phasor, the shadow of the sum is just the sum of the shadows, so we simply add the phasors as vectors, tip to tail, exactly as we added forces and displacements back in mechanics.

Here is the case we shall meet over and over. A resistor and an inductor carry the same current, so we lay the current along the horizontal reference. The voltage across the resistor, V_R, is in step with the current, so its phasor lies along the reference. The voltage across the inductor, V_L, leads the current by a quarter-cycle, so its phasor stands straight up at right angles. The supply voltage is the sum of the two, found by completing the right-angled triangle in the figure below.

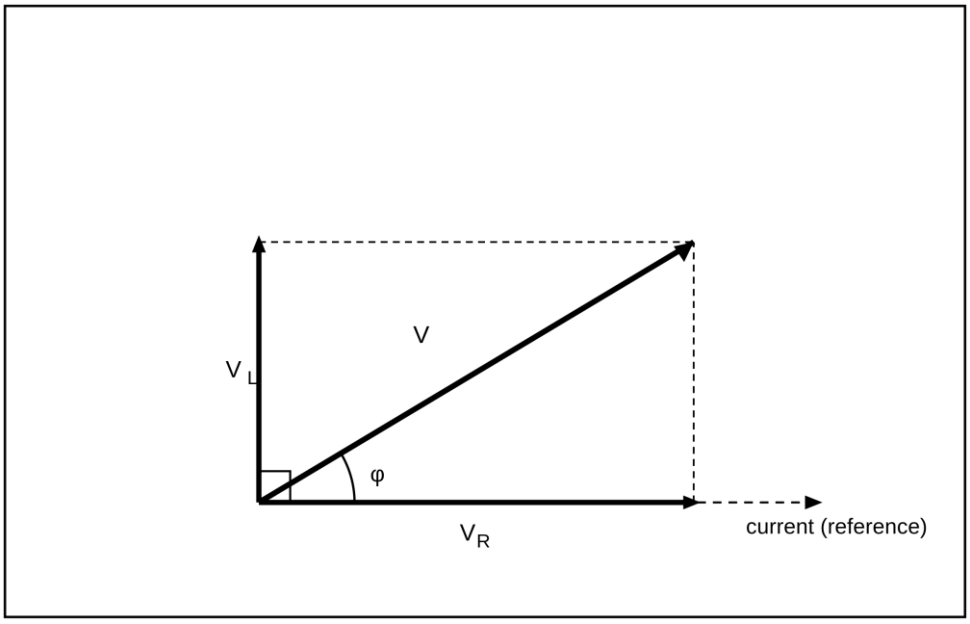


Figure: ***Adding two voltage phasors.*** *The resistor voltage lies along the current; the inductor voltage leads it by a quarter-cycle and points straight up. Their vector sum is the supply voltage V, the hypotenuse, which leads the current by the angle φ.*

Reading the triangle straight off, the peak supply voltage is the hypotenuse,

$$V = \sqrt{{V_R}^2 + {V_L}^2}$$

and it leads the current by the angle φ, given by:

$$\tan\varphi = \frac{V_L}{V_R}$$

Two arrows, one triangle, and Pythagoras: that is the entire method, and it will not change however many components we add.

Four examples to make the arrows behave. Kipanga, who once mistrusted ω, has decided he rather likes phasors, on the grounds that they are merely arrows and he has been drawing arrows since primary school. Let us not disabuse him.

BINDER Example 10

A voltage is represented by a phasor of length 50V rotating anticlockwise at angular frequency ω. Find the instantaneous voltage at the moment the phasor has turned through an angle of (a) 30°, (b) 90°, and (c) 210°.

Solution

The instantaneous value is the vertical projection of the phasor, $v = 50\sin\theta\ \text{V}$, where θ is the angle turned. So:

(a) $v = 50\sin 30°\text{V} = 50 \times 0.5\text{V} = 25\text{V}$

(b) $v = 50\sin 90°\text{V} = 50 \times 1\text{V} = 50\text{V}$

(c) $v = 50\sin 210°\text{V} = 50 \times (-0.5)\text{V} = -25\text{V}$

Making Sense of the Answer: *At* $90°$ *the phasor points straight up, so its projection is the full* $50V$ *peak. At* $210°$ *it has swung below the centre line, so the value has gone negative. Both readings match the picture exactly.*

Think Like a Physicist: *The projection is the instantaneous value. Reading a phasor is nothing more than dropping a perpendicular and measuring a height; no extra formula is needed.*

REAL Example 11

During a particularly slow afternoon lesson, Kipanga finds himself staring at the second hand of the classroom clock. Spotting an opportunity, Mr. Akilikubwa asks him to explain in what sense the second hand is a phasor. Help Kipanga answer by identifying what plays the role of (a) the phasor's length, (b) its angular frequency, and (c) its projection, and by stating the period and angular frequency of this particular phasor.

Solution

(a) The second hand is an arrow of fixed length, pinned at one end and turning steadily, which is precisely a phasor. Its length stands for the amplitude of the quantity it represents, here the greatest height its tip reaches above the centre of the clock.

(b) Its angular frequency is its steady rate of turning. The hand completes one full turn in 60s, so its period is $T = 60\text{s}$ and its angular frequency is:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{60\text{s}} = 0.105\text{rad s}^{-1}$$

(It happens to turn clockwise rather than anticlockwise, but that is merely a choice of direction; the physics of the projection is identical.)

(c) Its projection, the height of the tip above the centre line, rises and falls as a sine wave with that same period of 60s. If you filmed only the tip's height and plotted it against time, you would draw a perfect sinusoid, one full cycle every minute.

Making Sense of the Answer: *Every steadily rotating arrow is a phasor. The only thing special about a.c. is that the arrow stands for a voltage or current and spins far faster, fifty turns a second rather than one a minute.*

Think Like a Physicist: *When a new idea feels strange, find it inside something you already know. A phasor is just a clock hand whose shadow we happen to care about.*

HOT Example 12

In a series circuit the peak voltage across the resistor is 8.0V, in phase with the current, while the peak voltage across the inductor is 6.0V, leading the current by a quarter-cycle. Using a phasor diagram, find (a) the peak supply voltage and (b) the angle by which it leads the current.

Solution

(a) The two voltage phasors ($V_R = 8.0\text{V}, V_L = 6.0\text{V}$) meet at right angles, so the supply voltage is the hypotenuse of their triangle:

$$V = \sqrt{V_R^2 + V_L^2} = \sqrt{(8.0\text{V})^2 + (6.0\text{V})^2} = 10\text{V}$$

(b) The supply voltage leads the current by the angle φ, where:

$$\tan\varphi = \frac{V_L}{V_R} = \frac{6.0\text{V}}{8.0\text{V}} = 0.75, \qquad \text{so } \varphi = \tan^{-1}0.75 = 36.9°$$

Making Sense of the Answer: *The sides* 8, 6, 10 *are the familiar* 3-4-5 *triangle scaled up, a reassuring check. And a positive* φ *means the supply leads the current, exactly as it must when an inductor is in the circuit.*

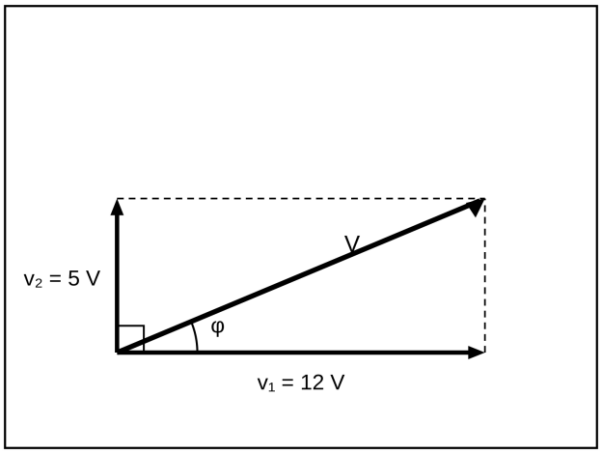
Think Like a Physicist: *Whenever two phasors meet at a right angle, the resultant is Pythagoras and the angle is an arctangent. You will lean on this one triangle more than any other relation in the chapter.*

HOT Example 13

Two alternating voltages of the same frequency act in series, $v_1 = 12\sin\omega t\ \text{V}$ and $v_2 = 5\sin(\omega t + 90°)\text{V}$. Find the single sinusoid that is their sum.

Solution

Draw each as a phasor, measuring angles from the horizontal. The first, v_1, has length 12V along the reference. The second, v_2, leads it by 90°, so its phasor of length 5V points straight up, as the following sketch shows.



The two phasors meet at a right angle, so the resultant is the hypotenuse,

$$\text{V} = \sqrt{(12\text{V})^2 + (5\text{V})^2} = 13\text{V}$$

at an angle ahead of the first phasor given by:

$$\tan\varphi = \frac{5\text{V}}{12\text{V}} = 0.417, \qquad \text{so } \varphi = \tan^{-1}0.417 = 22.6°$$

Since both phasors spin together at ω, the resultant spins with them, and its shadow is the sum we were after:

$$v = v_1 + v_2 = 13\sin(\omega t + 22.6°)\ \text{V}$$

Making Sense of the Answer: *The sides* 5, 12, 13 *form another whole-number right triangle, a clean check. The sum leads* v_1 *by only* 22.6°*, less than* v_2*'s full* 90° *lead, because* v_1 *is the larger arrow and pulls the resultant towards itself.*

Think Like a Physicist: *To add any number of same-frequency sinusoids, drop them all onto one phasor diagram, add the arrows as vectors, and read off the single resultant. Trigonometric identities are for those who have forgotten that they can draw.*

Four examples, and not one sine-addition formula in sight. Kipanga has noticed, with some satisfaction, that the frightening trigonometry of adding waves has quietly been replaced by the drawing of triangles, which he was good at all along.

PURE AC CIRCUITS: RESISTOR (R), INDUCTOR (L), CAPACITOR (C)

We now have the wave, a way to measure it, and the phasor to picture it. Time for the central question of the whole chapter. *When an alternating voltage is applied across each of the three basic components on its own, a resistor, an inductor, and a capacitor, how does the current answer?* Each answers differently, and those three answers are the alphabet from which every a.c. circuit is spelled. We take them one at a time, and for each we ask the same five questions: *is the current in step with the voltage, ahead of it, or behind? How strongly does the component oppose the current? Does it consume power? How does its behaviour change with frequency? And what happens at zero frequency, that is, on plain d.c.?* Master these three letters (R, L, C) and you can read any word.

First, the three circuits themselves. Each is the simplest thing imaginable, a single component connected straight across an alternating source:

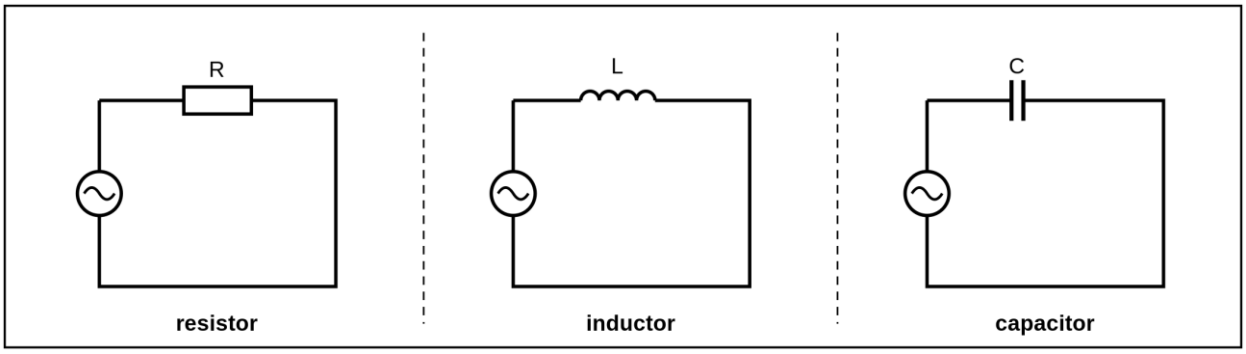


Figure: ***The three pure circuits of this section:*** *a resistor, an inductor, and a capacitor, each connected on its own across an alternating source. We ask the same five questions of each.*

And here, in a single glance, is the answer the section will earn, how the current sits relative to the voltage in each case:

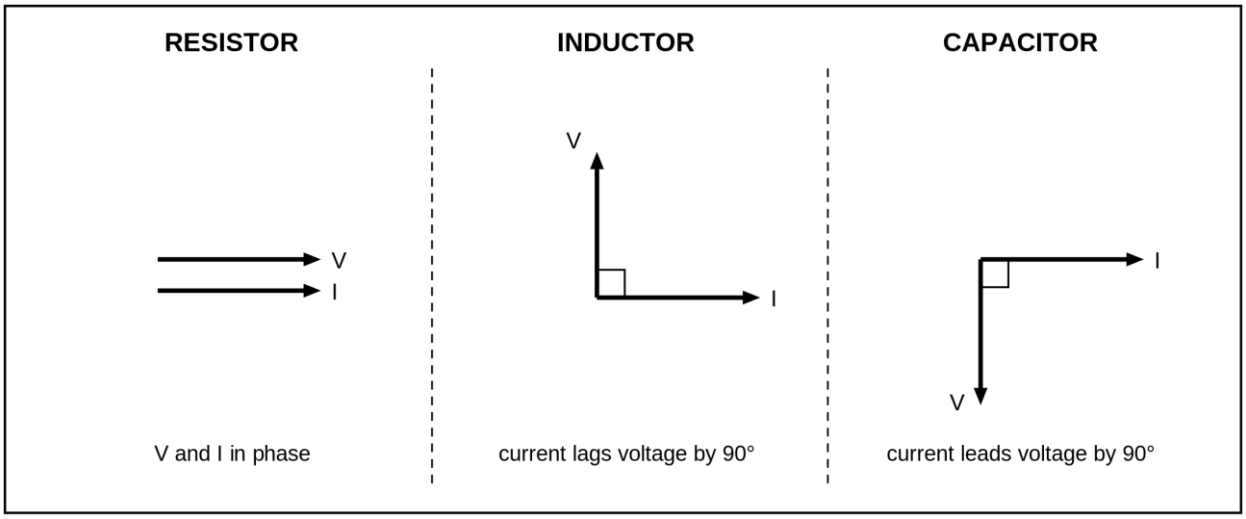


Figure: *The phase relationship between voltage and current in each of the three pure components, drawn as phasor diagrams with the current along the reference. In a resistor the two are in step; in an inductor the current lags the voltage by a quarter-cycle; in a capacitor the current leads the voltage by a quarter-cycle.*

The resistor: the honest baseline

Put an alternating voltage $v = V_0 \sin\omega t$ across a pure resistor and nothing surprising happens. Ohm's law from Chapter 12 still holds, instant by instant, $v = iR$, so the current is $i = \frac{v}{R} = \frac{V_0}{R}\sin\omega t = I_0 \sin\omega t$ with peak $I_0 = \frac{V_0}{R}$.

Run through the five questions.

- The current is exactly **in phase** with the voltage: they rise, peak, and cross zero together.
- The opposition is simply the **resistance R**, in ohms, the very same number as on d.c.; nothing new is needed.
- The resistor does **dissipate power**, and its average value is the one we built in the last section,

$$P_{av} = V_{rms} I_{rms} = I_{rms}^2 R$$

- Frequency makes no difference: R is R at any frequency, and the same at **zero frequency** too.

The resistor is the honest baseline against which the other two will look strange.

The inductor: the current lags behind

A pure inductor is a coil of wire with negligible resistance, and it makes its first appearance here, as a circuit component in its own right. A resistor is defined by Ohm's law, $v = iR$, and a capacitor by $i = C\frac{dv}{dt}$; an inductor is defined, in the same spirit, by its **inductance L**, measured in henries. Its defining rule is the exact mirror of the capacitor's: the voltage across the coil equals L times the rate at which the current through it changes:

$$\mathbf{v = L\frac{di}{dt}}$$

In plain words, *a coil does not care about the current itself, only about how fast that current is changing,* and it sets itself against that change. (Why a coil behaves this way is a matter for the deeper physics of magnetism, which we need not open up here.)

Watch what that rule forces. Suppose the current is $i = I_0 \sin\omega t$. Its rate of change is $\frac{di}{dt} = \omega I_0 \cos\omega t$ (differentiating a sine pulls down a factor of ω and turns it into a cosine, exactly as we saw for the spinning coil). The voltage across the inductor is therefore:

$$\mathbf{v = L\frac{di}{dt} = \omega L I_0 \cos\omega t}$$

Two results fall out of that one line. First, the voltage follows a cosine while the current follows a sine, and a cosine peaks a quarter-cycle before a sine, so **the current lags the voltage by 90°** (equivalently, the voltage leads the current). Second, the peak voltage is $V_0 = \omega L I_0$, so the ratio of peak voltage to peak current is:

$$\frac{V_0}{I_0} = \frac{\omega L I_0}{I_0} = \omega L$$

That ratio is the heart of the matter. For a resistor the ratio of voltage to current is the resistance R; for a coil it is ωL, and we give it its own name, the **inductive reactance $\mathbf{X_L}$**, measured in ohms exactly like a resistance:

$$X_L = \omega L = 2\pi f L$$

So **reactance is simply** *the a.c. opposition a component offers, the ratio of voltage to current*, and the current obeys an Ohm's-law lookalike,

$\mathbf{I_{rms} = \frac{V_{rms}}{X_L}}$ (the ratio is the same whether you use peak or rms values).

Two things, though, make reactance a genuinely new idea rather than a dressed-up resistance: it depends on frequency, and, as we are about to see, it costs no power. The frequency dependence is already in plain view, since ω sits right inside $X_L = \omega L$. Physically, the coil fights the current in proportion to how fast it is changing, and a higher frequency means a faster change at every instant, so the *reactance rises with frequency*. Double the frequency and you double the opposition.

Does it consume power? Not on average. During one quarter-cycle the source pushes the current up and pours energy into the inductor's magnetic field; during the next quarter-cycle the field collapses and hands every joule back to the source. Over a full cycle the **net power is zero**. This is the same bookkeeping as a swing in Chapter 6, where kinetic and potential energy traded back and forth without loss; here the trade is between the source and the magnetic field. Only a resistor actually spends energy.

Finally, zero frequency. On steady d.c. there is no change in the current for the inductor to fight, so $X_L = \omega L = 0$: the inductor becomes an ordinary piece of wire. An inductor passes d.c. freely and chokes a.c., the more harshly the higher the frequency, a fact we shall put to work when we meet the choke in a fluorescent lamp. A memory aid, offered only after the reasoning above: in an inductor the **E**MF leads the current (**I**), **ELI**.

The capacitor: the current leads

Now a pure capacitor. From Chapter 11 the charge it holds is $Q = CV$, and the current is the rate at which that charge flows on and off the plates, $i = C\frac{dv}{dt}$. Read that equation carefully: *the current depends not on the voltage itself but on how fast the voltage is changing.*

When the alternating voltage is changing fastest, as it sweeps through zero, charge floods on and off the plates and the current is greatest. When the voltage sits momentarily still at its peak, no charge moves and the current is zero. So the current reaches its maximum a quarter-cycle **before** the voltage does: **the current leads the voltage by a quarter-cycle, by 90°.** The capacitor is the mirror image of the inductor.

Put it in symbols, exactly as we did for the coil.

Let $v = V_0 \sin\omega t$.

Then the current is $i = C\frac{dv}{dt} = \omega C V_0 \cos\omega t$.

The current is a cosine to the voltage's sine, which confirms that **the current leads the voltage by 90°**; and the peak current is $I_0 = \omega C V_0$, so the ratio of peak voltage to peak current is:

$$\frac{V_0}{I_0} = \frac{V_0}{\omega C V_0} = \frac{1}{\omega C}$$

That ratio is the **capacitive reactance**, the a.c. opposition of a capacitor in ohms:

$$\mathbf{X_C = \frac{1}{\omega C} = \frac{1}{2\pi f C}}$$

with $I_{rms} = \frac{V_{rms}}{X_C}$ as before.

Notice the reciprocal, and with it the physical reason *the reactance falls with frequency*: the faster the voltage alternates, the less time the capacitor has to charge up before the polarity flips, so the less it manages to oppose, and the more current it passes. A capacitor finds high frequencies easy and low ones hard.

Power, again, averages to zero: energy stored in the electric field while the capacitor charges is returned in full while it discharges, so over a cycle the capacitor spends nothing. And at zero frequency the reciprocal bites hard: $X_C = \frac{1}{\omega C} = \infty$ when $\omega = 0$. A capacitor **blocks d.c. completely**: once charged, no further steady current flows. This is the very fact that bit Kipanga through the back of an unplugged television in Chapter 11, and the reason the capacitor of that chapter, fed a.c. instead, never settles but passes current cycle after cycle.

The matching memory aid: in a capacitor the current leads the EMF, **ICE**. Put the two together and generations of students have remembered the lot as ***ELI the ICE man***.

Reactance, frequency, and the d.c. limits

Step back and compare. Both X_L and X_C are measured in ohms and both play, for their component, the role resistance plays for a resistor. But unlike resistance they depend on frequency, and in opposite ways, as the graph below shows.

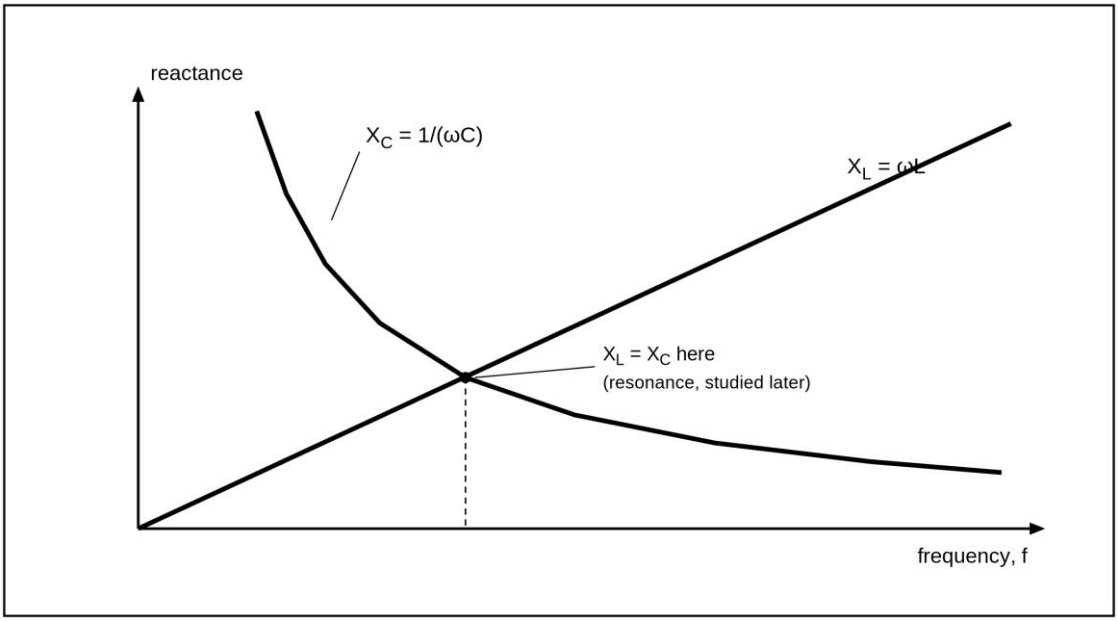


Figure: ***How the two reactances depend on frequency.*** *The inductive reactance grows in proportion to frequency, while the capacitive reactance falls as its reciprocal. At zero frequency (d.c.) the inductor offers no opposition while the capacitor offers infinite opposition; at very high frequency the roles reverse. The one frequency at which the two are equal is the seed of resonance, which we study later.*

Read the two limits straight off. At $f = 0$, the inductive reactance $X_L = \omega L = 0$ (a wire) while the capacitive reactance $X_C = \infty$ (a block); at very high frequency the inductor blocks while the capacitor waves the current through. An inductor and a capacitor are mirror images, and the single frequency where their reactances cross will turn out, in a later section, to be the most useful frequency in the whole circuit.

One last idea to plant. The amount of power a component draws depends on the angle φ between its voltage and its current, through $P_{av} = V_{rms} I_{rms} \cos\varphi$, a formula we develop properly in the power section. For a resistor $\varphi = 0$ and $\cos\varphi = \cos 0° = 1$: all the power is taken. For a pure inductor or capacitor $\varphi = 90°$ and $\cos\varphi = \cos 90° = 0$: none is. The quantity $\cos\varphi$, the **power factor**, will measure exactly how resistor-like or how reactance-like a real circuit behaves.

Five examples to drill the three letters. Kipanga, fresh from his victory over phasors, has promised to object the instant a current does something a current has no business doing. We are counting on him.

BINDER Example 14

A coil of inductance 0.50H and negligible resistance is connected to the 240V, 50Hz mains. Find (a) its inductive reactance and (b) the rms current it draws.

Solution

(a) The inductive reactance is:

$$X_L = \omega L = 2\pi f L = 2\pi \times 50\text{Hz} \times 0.50\text{H} = 157\Omega$$

(b) The rms current follows from the a.c. Ohm's law,

$$I_{rms} = \frac{V_{rms}}{X_L} = \frac{240\text{V}}{157\Omega} = 1.53\text{A}$$

Making Sense of the Answer: *A coil that would be almost a dead short on d.c. now opposes the current with* 157Ω*. That opposition is purely an a.c. effect, conjured entirely by the changing current.*

Think Like a Physicist: *For an inductor it is the reactance $X_L = \omega L$, not the resistance, that limits the a.c. current. Compute X_L first, then divide.*

BINDER Example 15

A 4.0μF capacitor is connected to the same 240V, 50Hz supply. Find (a) its capacitive reactance and (b) the rms current it draws.

Solution

(a) The capacitive reactance is:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50\text{Hz} \times 4.0 \times 10^{-6}\text{F}} = 796\Omega$$

(b) The rms current is:

$$I_{rms} = \frac{V_{rms}}{X_C} = \frac{240\text{V}}{796\Omega} = 0.302\text{A}$$

Making Sense of the Answer: *At $50Hz$ this small capacitor opposes the current more stiffly (796Ω) than the coil did (157Ω). Raise the frequency and that opposition would shrink.*

Think Like a Physicist: *Capacitive reactance is $X_C = \frac{1}{\omega C}$; mind the reciprocal. A smaller capacitor, or a lower frequency, means a larger reactance and a smaller current.*

REAL Example 16

Kipanga has been brooding over the last result, and finally he cannot hold it in.

Kipanga: *Sir, this one I refuse to accept. You say that in a capacitor the current leads the voltage, that the current reaches its peak a quarter-cycle before the voltage reaches its own. But the voltage is the cause and the current is the effect.* ***How can an effect arrive before its cause? Has the capacitor learned to see into the future?***

Mr. Akilikubwa: *A fair objection, and an honest one. But it hides a hidden assumption: that there was a first moment, a cause that started it all. In steady a.c. there is no such moment. The supply has been swinging back and forth for many cycles, and "leads" is not a story of cause and effect at all. It is a fixed relationship between two quantities that have been dancing together for a long time.*

The key is the capacitor's own law, $i = C\frac{dv}{dt}$. The current follows not the voltage but the slope of the voltage. The slope of a sine is steepest as the sine crosses zero and vanishes at the sine's peak. So the current is largest exactly when the voltage is changing fastest, and zero when the voltage is momentarily still. Trace those points and the current's peak falls a quarter-cycle ahead of the voltage's peak, not because it foresaw anything, but because the rate of change of a sine is itself a cosine, which simply peaks earlier. The lead is calculus, not prophecy.

Making Sense of the Answer: *Nothing is caused before its cause. The current at any instant is set by how fast the voltage is changing at that same instant, and that rate happens to peak a quarter-cycle before the voltage does.*

Think Like a Physicist: *When a "lead" or a "lag" feels paranormal, remember that a steady oscillation has no beginning. Phase is a standing relationship between quantities, never a claim that one event triggered another.*

HOT Example 17

A pure inductor of $L = 0.20\text{H}$ is connected across a 12V, 400Hz a.c. supply. Find (a) its reactance, (b) the rms current, and (c) the average power it consumes.

Solution

(a) The inductive reactance at this frequency is:

$$X_L = 2\pi fL = 2\pi \times 400\text{Hz} \times 0.20\text{H} = 503\Omega$$

(b) The rms current is:

$$I_{rms} = \frac{V_{rms}}{X_L} = \frac{12\text{V}}{503\Omega} = 0.0239\text{A} = 23.9\text{mA}$$

(c) The average power, with the phase angle $\varphi = 90°$ for a pure inductor, is:

$$P_{av} = V_{rms} I_{rms} \cos 90° = 0$$

Making Sense of the Answer: *A real current of $23.9mA$ flows, yet the inductor consumes no net power. The energy shuttles into the magnetic field and back out again each cycle; the source does work for half of every cycle and is repaid the other half.*

Think Like a Physicist: *Current flowing does not by itself mean power consumed. Only the part of the current in phase with the voltage carries power; a purely reactive current carries none.*

HOT Example 18

A 2.0μF capacitor is connected across a 12V supply. Find (a) the rms current when the supply is a.c. at 400Hz, (b) the average power it then consumes, and (c) the current that flows if the 12V supply is instead steady d.c.

Solution

(a) The capacitive reactance and current at 400Hz are:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 400\text{Hz} \times 2.0 \times 10^{-6}\text{F}} = 199\Omega$$

$$I_{rms} = \frac{V_{rms}}{X_C} = \frac{12\text{V}}{199\Omega} = 0.0603\text{A} = 60.3\text{mA}$$

(b) With $\varphi = 90°$ for a pure capacitor, $P_{av} = V_{rms} I_{rms} \cos 90° = 0$: no net power, for the same reason as the inductor, with the electric field now doing the storing.

(c) *On steady d.c. the frequency is zero,* so $X_C = \frac{1}{\omega C} = \infty$.

The capacitor charges up to 12V and then blocks all further flow: the steady current is **zero**.

Making Sense of the Answer: *The same capacitor passes* $60.3mA$ *of a.c. but not a whisker of d.c. The higher the frequency, the more it passes; at zero frequency it passes nothing. It is the exact mirror of the inductor, which passed d.c. freely and blocked the high frequencies.*

Think Like a Physicist: *To find what a capacitor does at any frequency, compute* $X_C = \frac{1}{\omega C}$. *Infinite at* $f = 0$ *(it blocks), small at high* f *(it passes). The reciprocal tells the whole story.*

Five examples, three letters, and one current caught apparently reading the future and then cleared of all charges. Kipanga objected, as promised, and was talked down; he now allows that a capacitor is no prophet, merely an unusually good judge of slopes, which he considers nearly as impressive.

SERIES COMBINATION CIRCUITS: RC, LR, AND RLC

In the last section we met the resistor, the inductor, and the capacitor each living alone. Real circuits string them together, and the simplest string is a **series** one, in which the same current flows through every component. We now learn to handle any series combination, and the phasor we built in the phasor section turns the job into something almost mechanical. The one new quantity we need is **impedance**, the a.c. cousin of resistance.

Impedance: the a.c. cousin of resistance

When resistors and reactances act together, *the total opposition the circuit offers to the alternating current* is called the **impedance**, written **Z** and measured in ohms. Like resistance, it is the ratio of voltage to current,

$$\mathbf{Z = \frac{V_{rms}}{I_{rms}}}$$

so the a.c. Ohm's law reads $V = IZ$. But impedance is more than a resistance. A **resistance** simply *dissipates*; a **reactance** simply *stores and returns*; **impedance** is the *combined effect of both*, and it *carries a phase angle as well as a size.* The crucial rule is that resistance and reactance are never added straight, because on the phasor diagram they sit at right angles to each other. They combine by Pythagoras, and that single fact,

$$Z = \sqrt{R^2 + X^2},$$

is the whole secret of every series circuit in this section.

A method that never changes

Before the particular circuits, here is the recipe they all obey. Learn it once and every series circuit, however it is dressed, yields to the same five steps.

Step 1: *Identify the components and the source: the rms (or peak) voltage, and the frequency.*

Step 2: *Find the reactances at that frequency,* $X_L = \omega L$ *and* $X_C = \frac{1}{\omega C}$.

Step 3: *Draw the phasor diagram with the current along the reference, build the impedance triangle, and read off* $Z = \sqrt{R^2 + (X_L - X_C)^2}$.

Step 4: *Apply the a.c. Ohm's law* $V_{rms} = I_{rms} Z$ *to find the unknown current or voltage.*

Step 5: *Find the phase angle from* $\tan\varphi = \frac{X_L - X_C}{R}$, *and state whether the current leads or lags.*

The three series circuits we shall study are drawn below. In each, the same current threads through every component in turn, driven round the loop by an alternating source.

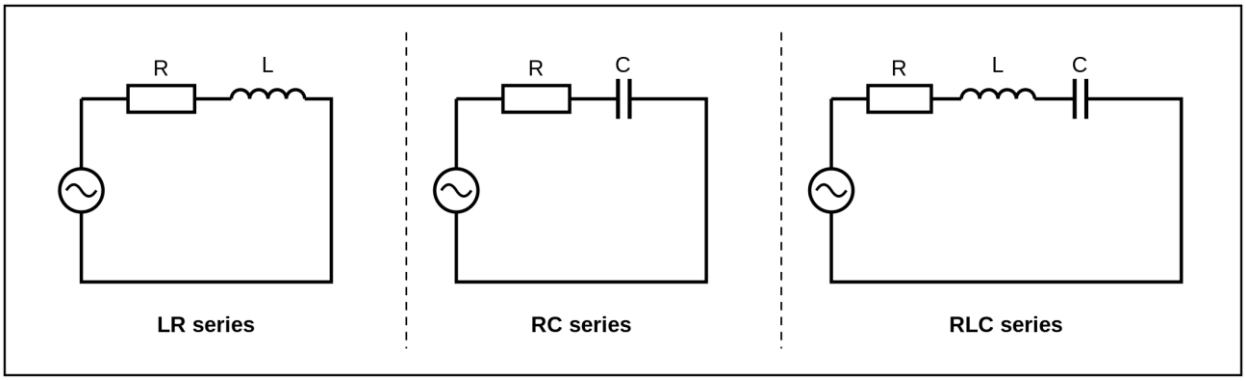


Figure: ***The three series circuits of this section:*** *a resistor with an inductor, a resistor with a capacitor, and all three together. In every case the same current flows through each component in turn, driven by an alternating source.*

The LR series circuit

Take a resistor and an inductor in series. The same current I flows through both, so we draw I along the reference. The resistor's voltage $V_R = IR$ is in phase with the current and lies along it; the inductor's voltage $V_L = IX_L$ leads the current by 90° and points straight up. The supply voltage is their phasor sum, the hypotenuse of the right-angled triangle in the figure below.

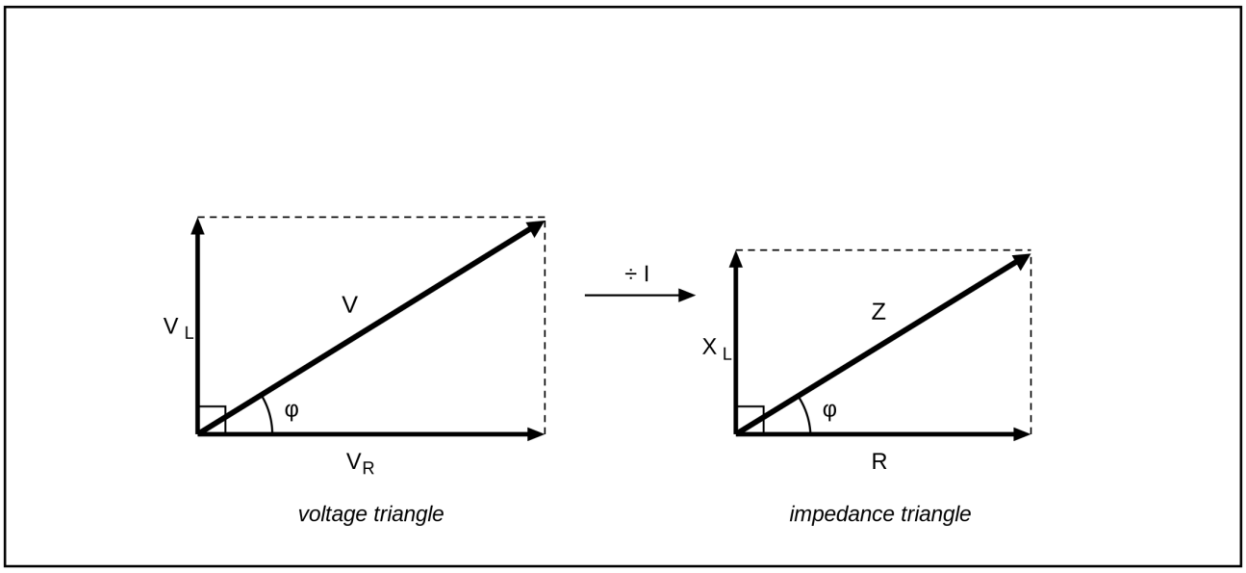


Figure: ***Adding the two voltages in an LR series circuit.*** *The resistor voltage is in phase with the current and the inductor voltage leads it by 90°, so the supply voltage V is the hypotenuse. Dividing every side by the common current I turns the voltage triangle into the impedance triangle, giving Z and the phase angle φ.*

Reading the triangle, the supply voltage is:

$$V = \sqrt{V_R^2 + V_L^2} = I\sqrt{R^2 + X_L^2}$$

Dividing through by the current gives the impedance,

$$\frac{V}{I} = Z = \sqrt{R^2 + X_L^2}$$

and the angle by which the supply voltage leads the current is given by:

$$\tan\varphi = \frac{X_L}{R} \text{ or } \varphi = \tan^{-1}\left(\frac{X_L}{R}\right)$$

The current **lags** the supply voltage; that is the mark of an inductive circuit.

This is exactly the circuit hiding inside a fluorescent tube. The tube of Chapter 12 will not run straight off the mains: left to itself the gas discharge draws more and more current until it destroys the tube. So a large inductor, the **choke** or ballast, is wired in series with it, making an LR circuit. Once the tube lights, the choke's reactance $X_L = 2\pi fL$ quietly limits the current to a safe value, dissipating almost no power of its own (a resistor large enough to do the same job would waste energy as heat). Every fluorescent tube humming in the Miono school corridor relies on this series inductor, or, in newer fittings, on an electronic circuit that imitates it.

The RC series circuit

A resistor and a capacitor in series is the mirror image. The same current flows; $V_R = IR$ lies along it, but the capacitor's voltage $V_C = IX_C$ now **lags** the current by 90°, so its phasor points straight down. The triangle is simply the LR one flipped below the current line, and the algebra follows at once:

$$Z = \sqrt{R^2 + X_C^2}, \qquad \tan\varphi = \frac{X_C}{R}$$

The only change of substance is the direction of the lean. Because the reactive voltage now points downward, the supply voltage lags the current, or, said the other way, **the current leads the supply voltage**. An inductor makes the current lag; a capacitor makes it lead.

The RLC series circuit

Now put all three in series. The same current flows through everything, so once more we draw it along the reference. The resistor's voltage lies along the current, the inductor's points up, and the capacitor's points down. Since V_L and V_C point in opposite directions, they partly cancel: the net reactive voltage is their difference $V_L - V_C$, as the figure below shows.

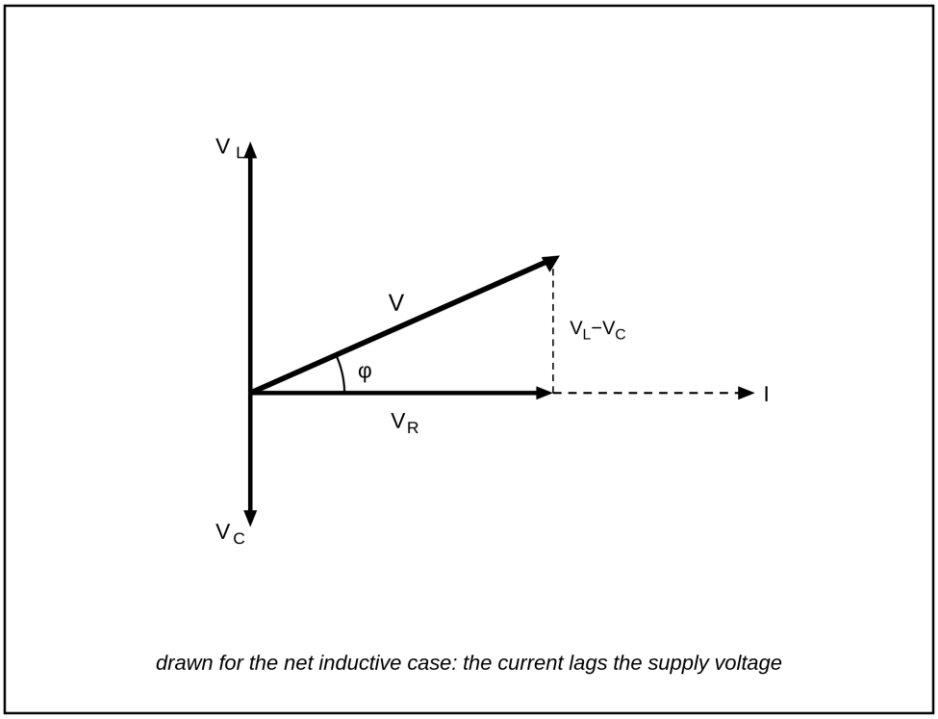


Figure: ***The phasor diagram of a series RLC circuit.*** *The inductor and capacitor voltages oppose each other, so only their difference survives, and the supply voltage V is the hypotenuse built on the resistor voltage and that net reactive voltage.*

The supply voltage is therefore the hypotenuse built on V_R and $(V_L - V_C)$, and dividing by the current gives the impedance and the phase angle:

$$Z = \sqrt{R^2 + (X_L - X_C)^2}, \qquad \tan\varphi = \frac{X_L - X_C}{R}$$

Everything depends on which reactance wins.

- When $X_L > X_C$, the circuit is net **inductive** and the current lags, behaving like an LR circuit.
- When $X_L < X_C$, it is net **capacitive** and the current leads, behaving like an RC circuit.

And when $X_L = X_C$, the two reactances cancel exactly, the impedance shrinks to its smallest possible value $Z = R$, and the current comes into step with the voltage. That last, special case is **resonance**, and it is important enough to have its own section later in the chapter.

Power factor

The phase angle φ measures how reactive a circuit is, and one reading of the impedance triangle captures it in a single number. Since R is the side adjacent to φ and Z is the hypotenuse,

$$\cos\varphi = \frac{R}{Z}$$

This ratio is the **power factor** we first named among the pure components. It runs from 1 for a purely resistive circuit ($\varphi = 0$, all the power is taken) down to 0 for a purely reactive one ($\varphi = 90°$, none is).

$$\mathbf{0 \leq \text{power factor} \leq 1}$$

For any series circuit, $\cos\varphi = \frac{R}{Z}$, and we shall see in the power section how directly it controls the bill a workshop pays.

Five examples to put the five steps to work. Kipanga, who has come to trust triangles, is now positively eager, and has begun drawing phasor diagrams.

BINDER Example 19

In an LR series circuit the resistance is 30Ω and the inductive reactance at the working frequency is 40Ω. Find: (a) the impedance and (b) the angle by which the current lags the supply voltage.

Solution

(a) Resistance and reactance combine by Pythagoras:

$$Z = \sqrt{R^2 + X_L^2} = \sqrt{(30\Omega)^2 + (40\Omega)^2} = 50\Omega$$

(b) The phase angle is

$$\tan\varphi = \frac{X_L}{R} = \frac{40\Omega}{30\Omega} = 1.33, \quad \varphi = \tan^{-1}1.33 = 53.1^\circ$$

Making Sense of the Answer: *The sides* $30, 40, 50$ *are the* 3-4-5 *triangle again. The impedance,* 50Ω*, is larger than either the resistance or the reactance alone, exactly as a hypotenuse must be.*

Think Like a Physicist: *Never add R and* X_L *straight; they are at right angles. Combine them with Pythagoras and read the phase from the arctangent of the ratio.*

REAL Example 20

The fluorescent tubes in the Miono school corridor each have a small, heavy component wired in series with the tube, called a choke. A new technician, replacing a dead tube, wonders whether he can save money by leaving the choke out, since "it is only a coil of wire and does nothing useful." Explain what the choke does and what would happen to the tube without it.

Solution

The choke is a large inductor, and the tube together with it forms an LR series circuit. The reason it is essential lies in the strange behaviour of the gas inside the tube: once the gas has been made to conduct, its resistance falls as more current flows, so a tube connected straight to the mains would draw an ever-larger current, overheat, and quickly destroy itself. The choke prevents this runaway. Its reactance $X_L = 2\pi fL$ rises with the current's demands and limits the current to a safe steady value, holding the tube at its proper operating point.

Why not use a plain resistor to limit the current instead? Because a resistor would dissipate a great deal of power as wasted heat, whereas the choke, being almost pure reactance, limits the current while consuming almost no power of its own. The technician's "useless coil" is in fact what keeps the tube alive and the electricity bill low. Leave it out and the new tube will flare brilliantly for a moment and then burn out.

Making Sense of the Answer: *This is the zero-power property of a pure inductor put to practical use: the choke controls the current without paying for it in heat, which a resistor never could.*

Think Like a Physicist: *A reactance is the engineer's tool of choice whenever a current must be limited without wasting energy. Resistors limit and waste; reactances limit and store.*

HOT Example 21

A $4.0\mu F$ capacitor is connected in series with an 800Ω lamp across a $9.0V$, $60Hz$ a.c. supply. Find (a) the capacitive reactance, (b) the impedance, (c) the rms current, (d) the voltage across the capacitor and across the lamp, (e) the phase angle, and (f) the average power dissipated.

Solution

(a) The capacitive reactance is:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 60\text{Hz} \times 4.0 \times 10^{-6}\text{F}} = 663\Omega$$

(b) The lamp behaves as a resistor, so the impedance is:

$$Z = \sqrt{R^2 + X_C^2} = \sqrt{(800\Omega)^2 + (663\Omega)^2} = 1039\Omega$$

(c) The rms current follows from the a.c. Ohm's law:

$$I_{rms} = \frac{V_{rms}}{Z} = \frac{9.0V}{1039\Omega} = 8.66\text{mA}$$

(d) The two component voltages are:

$$V_C = I_{rms}X_C = 8.66\text{mA} \times 663\Omega = 5.74V$$
$$V_R = I_{rms}R = 8.66\text{mA} \times 800\Omega = 6.93V$$

(e) The phase angle is:

$$\tan\varphi = \frac{X_C}{R} = \frac{663\Omega}{800\Omega} = 0.829, \quad \varphi = \tan^{-1}0.829 = 39.6°$$

with the current leading the supply voltage.
(f) The average power is dissipated only in the lamp:

$$P_{av} = I_{rms}^2 R = (8.66\text{mA})^2 \times 800\Omega = 0.060\text{W} = 60\text{mW}$$

Making Sense of the Answer: *A quick check:* $\sqrt{V_R^2 + V_C^2} = \sqrt{(6.93V)^2 + (5.74V)^2} = 9.0V$*, the supply voltage, as the phasor triangle demands. The capacitor and lamp voltages do not add to* $9.0V$ *straight, because they are a quarter-cycle apart.*

Think Like a Physicist: *The component voltages in a series a.c. circuit add as phasors, not as numbers. If* V_R *and* V_C *ever seem to "add up to more than the supply", remember they are at right angles, and Pythagoras restores order.*

HOT Example 22

A resistor $R = 200\Omega$, an inductor $L = 2.0\text{H}$, and a capacitor $C = 10\mu\text{F}$ are connected in series across a 240V, 50Hz supply. Find the two reactances, the impedance, the rms current, the phase angle, and the average power.

Solution
The two reactances at 50Hz are:

$$X_L = 2\pi fL = 2\pi \times 50\text{Hz} \times 2.0\text{H} = 628\Omega$$

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50\text{Hz} \times 10 \times 10^{-6}\text{F}} = 318\Omega$$

The net reactance is $X_L - X_C = 628\Omega - 318\Omega = 310\Omega$, inductive (since $X_L > X_C$). The impedance is:

$$Z = \sqrt{R^2 + (X_L - X_C)^2} = \sqrt{(200\Omega)^2 + (310\Omega)^2} = 369\Omega$$

The rms current and phase angle are then:

$$I_{rms} = \frac{V_{rms}}{Z} = \frac{240\text{V}}{369\Omega} = 0.650\text{A}$$

$$\tan\varphi = \frac{X_L - X_C}{R} = \frac{310\Omega}{200\Omega} = 1.55, \quad \varphi = \tan^{-1}1.55 = 57.2°$$

and, since the circuit is net inductive, the current lags the supply voltage by $57.2°$.
The average power is dissipated only in the resistor:

$$P_{av} = I_{rms}^2 R = (0.650\text{A})^2 \times 200\Omega = 84.5\text{W}$$

Making Sense of the Answer: *The inductor's* 628Ω *and the capacitor's* 318Ω *fight each other, and only the* 310Ω *that survives joins the resistance. Had the two reactances been equal, the impedance would have collapsed to the bare* 200Ω *of the resistor.*

Think Like a Physicist: *In an RLC circuit always subtract the reactances first,* $X_L - X_C$*, before reaching for Pythagoras. The sign of that difference tells you at a glance whether the current will lead or lag.*

HOT Example 23

A real coil possesses both resistance and inductance. It draws 0.50A from a 12V d.c. supply, and 0.24A from a 12V, 50Hz a.c. supply. Find: (a) the coil's resistance, (b) its impedance on a.c., (c) its inductive reactance, and (d) its inductance.

Solution
(a) On d.c. the frequency is zero, so the reactance $X_L = \omega L = 0$ and only the resistance opposes the current. Hence:

$$R = \frac{12\text{V}}{I_{dc}} = \frac{12\text{V}}{0.50\text{A}} = 24\Omega$$

(b) On a.c. the full impedance opposes the current:

$$Z = \frac{12\text{V}}{I_{ac}} = \frac{12\text{V}}{0.24\text{A}} = 50\Omega$$

(c) Since $Z = \sqrt{R^2 + X_L^2}$, the reactance is:

$$X_L = \sqrt{Z^2 - R^2} = \sqrt{(50\Omega)^2 - (24\Omega)^2} = 43.9\Omega$$

(d) Finally the inductance, from $X_L = 2\pi fL$:

$$L = \frac{X_L}{2\pi f} = \frac{43.9\Omega}{2\pi \times 50\text{Hz}} = 0.14\text{H}$$

Making Sense of the Answer: *The trick is that d.c. sees only* R *while a.c. sees the whole of* Z*. Two simple measurements, one battery and one a.c. source, are enough to separate a coil's resistance from its inductance.*

Think Like a Physicist: *Whenever a component is tested on both d.c. and a.c., let the d.c. reading hand you* R *and the a.c. reading hand you* Z*; the reactance is then* $\sqrt{Z^2 - R^2}$*.*

Five circuits, one method, and not a moment's panic. Kipanga has noticed that every example ended in the same triangle, and has concluded, not unreasonably, that the whole of a.c. is Pythagoras wearing a disguise.

PARALLEL COMBINATION CIRCUITS: LR, CR, AND RLC

The series circuits of the last section made the components stand in single file, so that one and the same current had to pass through every one of them in turn. A **parallel** circuit does the opposite. It connects each component straight across the supply, between the very same two points, so that now it is the **voltage** that is shared by all of them and the current that is free to split. Imagine a daladala stand in Dar es Salaam: in a series circuit every passenger files aboard through one door, one after another; in a parallel circuit each climbs in by a door of its own, and only the conductor, counting heads, knows the total. Everything we learned about phasors still holds, but two of the rules quietly trade places.

In a series circuit the current was common and the voltages added as phasors. In a parallel circuit the **voltage is common** and it is the **currents** that add as phasors. So this time we draw the voltage as the reference phasor, work out the current in each branch on its own, and add those branch currents head to tail to find the total current drawn from the supply. That total is called the **line current**.

Here are a resistor, an inductor, and a capacitor all hung in parallel across one a.c. supply.

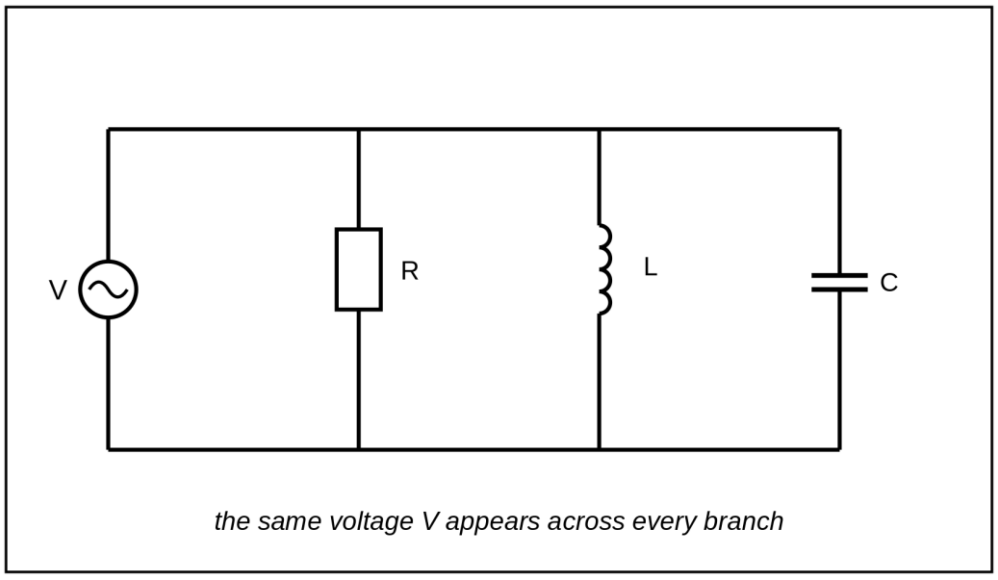


Figure: ***A resistor, an inductor, and a capacitor connected in parallel across one a.c. supply.*** *The same voltage drives all three branches, and the current the supply delivers is the sum of the three branch currents.*

One voltage, three branch currents

Because the same voltage sits across every branch, we may treat each branch entirely on its own, exactly as in the pure-circuit section, and simply read off its current. Let V be the rms supply voltage. Then:

$$I_R = \frac{V}{R}, \quad I_L = \frac{V}{X_L}, \quad I_C = \frac{V}{X_C}$$

and each branch current keeps the phase rule we proved for it earlier: the resistor's current stays **in step** with the voltage, the inductor's **lags** a quarter cycle behind, and the capacitor's **leads** a quarter cycle ahead. By Kirchhoff's current law from Chapter 12, the current the supply delivers is the sum of these three branch currents, but because they are out of step with one another we must add them as **phasors**, not as plain numbers.

Adding the branch currents

Lay the supply voltage V along the reference direction. The resistor current I_R lies right along it. The inductor current I_L hangs a quarter turn behind, drawn downwards, and the capacitor current I_C stands a quarter turn ahead, drawn upwards. The two reactive currents point in exactly opposite directions, so they fight each other, and only their difference survives.

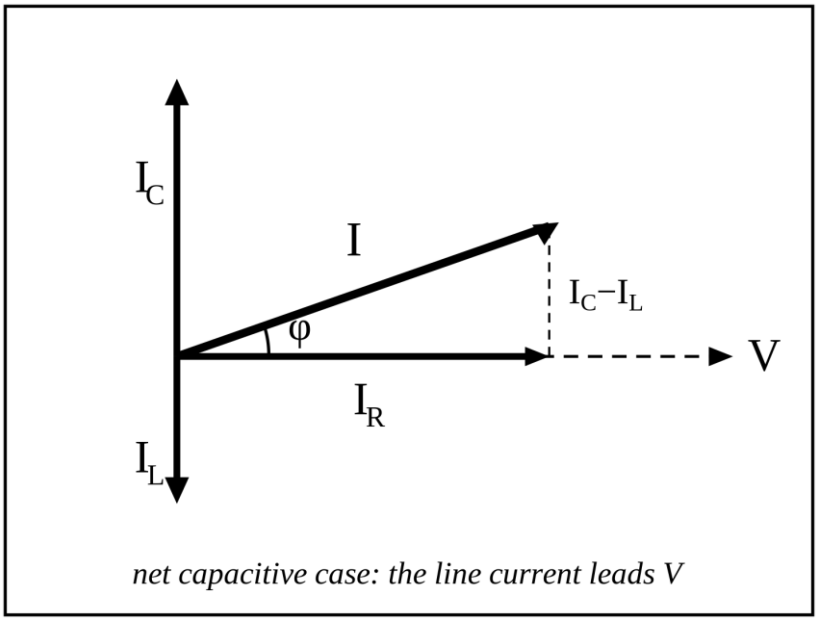


Figure: ***The branch currents drawn as phasors, with the common voltage along the reference.*** *The resistor current lies along the voltage, the capacitor current leads by a quarter cycle, and the inductor current lags by a quarter cycle. Their head-to-tail sum is the line current.*

Adding them head to tail gives a right-angled triangle once more, this time with the resistor current along the base and the net reactive current up the side. The line current is therefore:

$$I = \sqrt{I_R^2 + (I_C - I_L)^2}$$

and it makes an angle φ with the supply voltage given by:

$$\tan\varphi = \frac{I_C - I_L}{I_R},$$

leading if I_C is the larger (net capacitive) and lagging if I_L is the larger (net inductive).

Notice how the picture has turned inside out compared with the series circuit. There it was V_L that pointed up and V_C down; here it is I_C that points up and I_L down. Voltage and current have swapped their geometry, because in going from series to parallel the reactance and its reciprocal have swapped their roles.

Impedance, and a shortcut called admittance

The impedance of the whole combination is still defined the old way, as the ratio of the supply voltage to the current it drives:

$$Z = \frac{V}{I}$$

So once the line current is known the impedance follows in a single step. For parallel work there is also a neater quantity waiting in the wings. *The reciprocal of impedance* is called the **admittance**, written **Y** and measured in **siemens** (the symbol is **S**):

$$\mathbf{Y = \frac{1}{Z} = \frac{I}{V}}$$

Admittance measures how **easily** current flows rather than how strongly it is opposed, and in a parallel circuit the admittances of the separate branches simply add, which is often the quickest route of all. We shall keep impedance as our main tool and reach for admittance only when it plainly saves work.

A current that almost disappears

Look once more at the phasor triangle and try a thought experiment. The inductor current points down, the capacitor current points up. Suppose we tune the frequency until the two reactances become equal, $X_L = X_C$. Then the two branch currents are equal as well, $I_L = I_C$, and being equal and opposite they cancel completely. The net reactive current vanishes, the line current shrinks to the bare resistor current, and the impedance $Z = \frac{V}{I}$ swells towards its **largest** value.

This is the exact opposite of the series circuit, where equal reactances made the impedance collapse to a minimum and the current roar to a maximum. A series circuit at that special frequency is greedy for current; a parallel one is miserly with it. That contrast is the doorway to **resonance**, and we step through it in the next section.

Handling any parallel circuit, in five steps

Step 1: *Draw it, mark the shared voltage. The same voltage V sits across every branch; that is your reference phasor.*

Step 2: *Reactances, then branch currents. Find $X_L = \omega L$ and $X_C = \frac{1}{\omega C}$, then $I_R = \frac{V}{R}$, $I_L = \frac{V}{X_L}$, $I_C = \frac{V}{X_C}$.*

Step 3: *Place the phasors. I_R along V; I_L a quarter turn behind; I_C a quarter turn ahead.*

Step 4: *Combine. Net reactive current $I_C - I_L$, then $I = \sqrt{I_R^2 + (I_C - I_L)^2}$.*

Step 5: *Angle and impedance. $tan\varphi = \frac{I_C - I_L}{I_R}$; then $Z = \frac{V}{I}$ if it is wanted.*

Time to send some current down these branches. Kipute has been promised that parallel circuits are merely "series circuits standing on their heads," and is braced for the worst. Four circuits lie ahead, one shared voltage apiece, and at least one supply current that comes out **smaller** than the branch currents it is feeding. Mr. Akilikubwa calls that last one "the miracle of the loaves and the currents" and refuses to explain the trick until the arithmetic has been done.

BINDER Example 24

A resistor of 60Ω and a pure inductor of reactance 80Ω are connected in parallel across a 240V a.c. supply. Find: (a) the current in each branch, (b) the line current drawn from the supply, and (c) the phase angle between the line current and the supply voltage.

Solution

(a) The supply voltage sits across both branches. The resistor current is in phase with it, the inductor current a quarter cycle behind:

$$\text{I}_\text{R} = \frac{240\text{V}}{60\Omega} = 4\text{A}, \quad \text{I}_\text{L} = \frac{240\text{V}}{80\Omega} = 3\text{A}$$

(b) The two branch currents are $90°$ apart, so they add by Pythagoras:

$$\text{I} = \sqrt{\text{I}_\text{R}^2 + \text{I}_\text{L}^2} = \sqrt{(4\text{A})^2 + (3\text{A})^2} = 5\text{A}$$

(c) The phase angle follows from

$$\tan\varphi = \frac{\text{I}_\text{L}}{\text{I}_\text{R}} = \frac{3\text{A}}{4\text{A}} = 0.75, \quad \varphi = \tan^{-1}0.75 = 36.9°$$

and, the only reactive branch being the inductor, the line current **lags** the supply voltage by $36.9°$.

Making Sense of the Answer: *The familiar 3, 4, 5 triangle, but built this time from currents rather than voltages. The line current $(5A)$ is larger than either branch current, because the resistor and the inductor each take their share and the supply must feed both at once.*

Think Like a Physicist: *In a parallel circuit you add currents; in a series circuit you add voltages. Same Pythagoras, different quantity. The first question to ask of any circuit is simply: which quantity is the shared one?*

BINDER Example 25

A resistor of 40Ω and a capacitor of reactance 30Ω are joined in parallel across a 120V a.c. supply. Find the line current and its phase relative to the supply voltage.

Solution

The branch currents are

$$\text{I}_\text{R} = \frac{120\text{V}}{40\Omega} = 3\text{A}, \quad \text{I}_\text{C} = \frac{120\text{V}}{30\Omega} = 4\text{A}$$

The capacitor current leads by $90°$, so it stands at right angles to the resistor current. The line current is:

$$\text{I} = \sqrt{\text{I}_\text{R}^2 + \text{I}_\text{C}^2} = \sqrt{(3\text{A})^2 + (4\text{A})^2} = 5\text{A}$$

and the phase angle is:

$$\tan\varphi = \frac{\text{I}_\text{C}}{\text{I}_\text{R}} = \frac{4\text{A}}{3\text{A}} = 1.33, \quad \varphi = \tan^{-1}1.33 = 53.1°$$

This time the reactive branch is capacitive, so the line current **leads** the supply voltage by $53.1°$.

Making Sense of the Answer: *The mirror image of the previous example. Swap the inductor for a capacitor and the only thing that changes is the direction of the lean: the line current now runs ahead of the voltage instead of behind it.*

Think Like a Physicist: *Lead or lag is settled by which reactive branch is present. Capacitor in the lead, inductor in the rear. The sizes follow the same triangle either way.*

HOT Example 26

An inductor of reactance 40Ω and a capacitor of reactance 50Ω, each of negligible resistance, are connected in parallel across a $200V$ supply. Find the current in each branch and the current drawn from the supply, and comment on the result.

Solution

Each branch carries:

$$I_L = \frac{200V}{40\Omega} = 5A, \quad I_C = \frac{200V}{50\Omega} = 4A$$

The inductor current lags the voltage by $90°$ and the capacitor current leads it by $90°$, so the two are a full $180°$ apart, in direct opposition. With no resistor branch there is no in-phase current at all, and the line current is simply their difference:

$$I = I_L - I_C = 5A - 4A = 1A$$

Because the inductor current is the larger, the surviving line current lags the supply voltage by $90°$. The impedance of the combination is:

$$Z = \frac{V}{I} = \frac{200V}{1A} = 200\Omega$$

Comment: The supply delivers a mere $1A$, and yet $5A$ surges through the inductor and $4A$ through the capacitor.

Making Sense of the Answer: *The branch currents dwarf the line current. A hidden $4A$ simply sloshes back and forth between capacitor and inductor every cycle, the capacitor handing its stored charge to the coil and the coil handing it straight back, while the supply tops up only the $1A$ of difference. The closer the two reactances, the smaller the line current and the larger the impedance, heading towards the parallel resonance of the next section.*

Think Like a Physicist: *When two reactive branches face each other with no resistor between them, subtract the currents, do not Pythagoras them; they are $180°$ apart, not $90°$. A parallel inductor and capacitor can together oppose the supply far more stubbornly than either does alone.*

HOT Example 27

A resistor of 100Ω, an inductor of $0.3H$, and a capacitor of $20\mu F$ are connected in parallel across a $240V$, $50Hz$ supply. Find (a) the three branch currents, (b) the line current, (c) the phase angle, and (d) the impedance of the combination.

Solution

(a) First the two reactances, then the three branch currents. The angular frequency is:

$$\omega = 2\pi f = 2\pi \times 50Hz = 314rad/s$$

It follows that:

$$X_L = \omega L = 314rad/s \times 0.3H = 94.2\Omega$$

$$X_C = \frac{1}{\omega C} = \frac{1}{314rad/s \times 20 \times 10^{-6}F} = 159\Omega$$

So the branch currents are:

$$I_R = \frac{240V}{100\Omega} = 2.40A, \quad I_L = \frac{240V}{94.2\Omega} = 2.55A, \quad I_C = \frac{240V}{159\Omega} = 1.51A$$

(b) The resistor current is in phase; the reactive currents oppose, leaving a net reactive current $I_C - I_L = 1.51A - 2.55A = -1.04A$ (negative, hence net inductive). The line current is:

$$I = \sqrt{{I_R}^2 + (I_C - I_L)^2} = \sqrt{(2.40A)^2 + (1.04A)^2} = 2.62A$$

(c) The phase angle:

$$\tan\varphi = \frac{I_C - I_L}{I_R} = \frac{-1.04A}{2.40A} = -0.433, \quad \varphi = 23.4°$$

The minus sign tells us the line current **lags** the supply voltage by $23.4°$, the inductive branch having won.

(d) The impedance of the whole combination is:

$$Z = \frac{V}{I} = \frac{240V}{2.62A} = 91.6\Omega$$

Making Sense of the Answer: *The inductor draws the most current ($2.55A$) because at $50Hz$ its reactance is the smallest; the capacitor draws the least. Their tug of war leaves the circuit slightly inductive, so the line current lags, exactly as the negative net reactive current warned.*

Think Like a Physicist: *In a parallel RLC, always test the sign of $I_C - I_L$. Positive means net capacitive and a leading line current; negative means net inductive and a lagging one. The resistor current sets the in-phase backbone that the reactive difference leans against.*

Four circuits dispatched, and one supply caught feeding a single amp into a pair of branches happily trading five and four amps between themselves. Kipanga, who began the section certain that the total must always be the biggest number in sight,

has gone very quiet. Kipute, sensing weakness, has taken to calling the parallel inductor-and-capacitor "the circuit that prints its own current," which is not quite right, but is the sort of wrong that means he has understood it.

RESONANCE: THE CIRCUIT THAT CHOOSES A FREQUENCY

All through the series section we kept subtracting one reactance from the other, $X_L - X_C$, and all through the parallel section we watched the two reactive currents fight. A restless student eventually asks the obvious question: *what if the two were exactly equal? What if the tug of war ended in a perfect draw?* The answer turns out to be the single most useful idea in the whole of a.c. theory. It is the reason a radio can pluck one station out of the crowded air and ignore all the rest, and it goes by the name **resonance**.

You have in fact met resonance before, in a quieter form. A child on a swing, pushed gently but in time with the swing's own rhythm, rises higher and higher; pushed out of time, she barely moves. Every object that can oscillate has a natural frequency it "prefers," and when you drive it at that frequency the response grows large. We studied this for mechanical oscillators back in the chapter on simple harmonic motion. A circuit holding both an inductor and a capacitor has exactly such a natural frequency, and shows exactly the same dramatic response when driven at it.

The frequency where the reactances cancel

In a series RLC circuit the impedance is $Z = \sqrt{R^2 + (X_L - X_C)^2}$. The resistance is fixed, but the two reactances change with frequency in opposite ways: $X_L = \omega L$ grows as the frequency rises, while $X_C = \frac{1}{\omega C}$ shrinks. At low frequency the capacitor dominates; at high frequency the inductor dominates. Somewhere in between there must be one frequency at which the two are exactly equal and cancel. There the bracket vanishes, the impedance falls to its smallest possible value, just R, and the circuit behaves as though the inductor and capacitor were not present at all.

To find that frequency, set the two reactances equal:

$$\omega L = \frac{1}{\omega C}$$

Multiplying both sides by ω and dividing by L gives:

$$\omega^2 = \frac{1}{LC}, \quad \text{so} \quad \omega_0 = \frac{1}{\sqrt{LC}}$$

and since $\omega = 2\pi f$, the **resonant frequency** is:

$$\mathbf{f_0 = \frac{1}{2\pi\sqrt{LC}}}$$

This is the circuit's own natural frequency, fixed entirely by L and C; the resistance has no say in it at all. *The frequency at which a circuit's reactances cancel and its response is greatest* is called its **resonant frequency**, and the circuit is said to be in **resonance**.

What happens at resonance in a series circuit

Imagine the series circuit driven exactly at f_0.

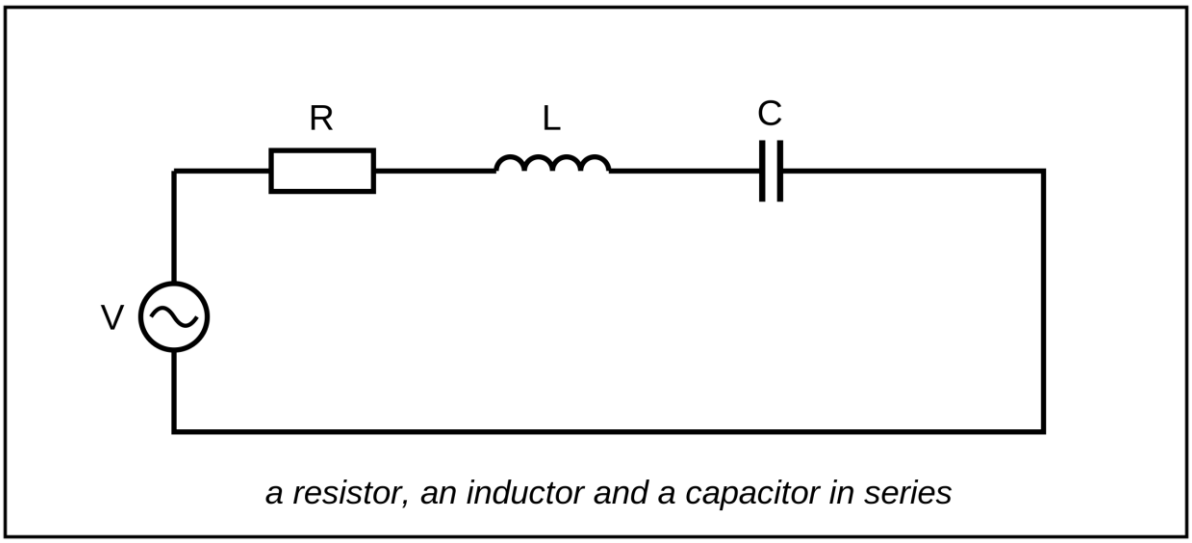


Figure: ***A resistor, an inductor and a capacitor in series across an alternating supply.*** *When the supply frequency equals the resonant frequency, the inductor's and capacitor's reactances cancel and the circuit offers only its resistance.*

Four things happen together, and each is worth stating plainly.

First, the impedance is at its minimum, $Z = R$, purely resistive.

Second, the current is at its maximum, $I = \frac{V}{R}$, limited only by the resistance.

Third, the current is exactly in step with the supply voltage; the phase angle is zero, and (as the power section will show) the circuit then draws energy as efficiently as it possibly can.

Fourth, and most surprising of all, the voltages across the inductor and the capacitor are equal in size and opposite in direction, so they cancel within the loop, and yet each one on its own can be many times larger than the supply voltage itself.

If we plot the current against the supply frequency, the result is the famous **resonance curve**: a peak that climbs to its maximum at f_0 and falls away on either side. The impedance does the mirror image, dipping to its minimum value R at f_0 and rising on both sides. A lamp placed in the circuit would glow brightest exactly at resonance and dim away as the frequency moved off it in either direction.

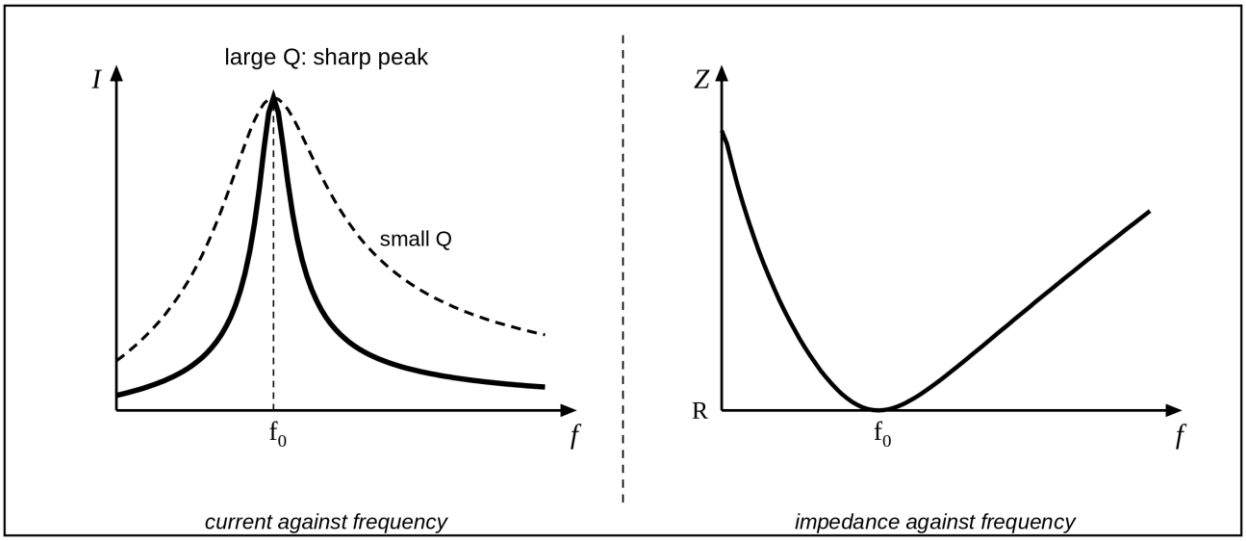


Figure: ***Left: the current rises to a sharp peak at the resonant frequency.*** *A circuit with a large Q gives a tall, narrow peak (solid); a small Q gives a low, broad one (dashed). Right: the impedance dips to its minimum value R at the same frequency.*

Sharpness, the Q-factor, and voltage magnification

How sharp is the peak? That is measured by the **quality factor**, or **Q-factor**, of the circuit:

$$\mathbf{Q = \frac{\omega_0 L}{R} = \frac{1}{\omega_0 CR} = \frac{1}{R}\sqrt{\frac{L}{C}}}$$

A large Q means a tall, narrow peak: the circuit responds strongly, but only to a very narrow band of frequencies around f_0. A small Q gives a low, broad peak. The width of the peak is called the **bandwidth**, and it is:

$$\mathbf{\Delta f = \frac{f_0}{Q}}$$

The Q-factor carries a second meaning, just as important. At resonance the voltage across the inductor (and the equal, opposite voltage across the capacitor) is Q times the supply voltage:

$$V_L = V_C = QV$$

So a circuit with $Q = 100$ driven by a 1V supply develops a full 100V across its capacitor. This **voltage magnification** is real, occasionally dangerous, and exactly what makes a tuned circuit so sensitive.

This sharpness is the secret of tuning. The aerial of a radio picks up the signals of every station at once. Behind it sits a circuit whose capacitor can be varied; turning the tuning knob changes C and so changes f_0. When f_0 is brought to match the frequency of one station, that station drives the circuit at resonance and is magnified, while all the others, off resonance, are passed over. The higher the Q, the more sharply the set tells one station from its neighbour.

Parallel resonance: the rejector circuit

In the parallel section we noticed something and promised to come back to it. When an inductor and a capacitor are placed in parallel and the frequency is tuned so that $X_L = X_C$, the two branch currents become equal and opposite and cancel, the line current drops almost to nothing, and the impedance of the combination soars to a maximum. That is **parallel resonance**, and the inductor-capacitor pair is called a **tank circuit**, or a **rejector**.

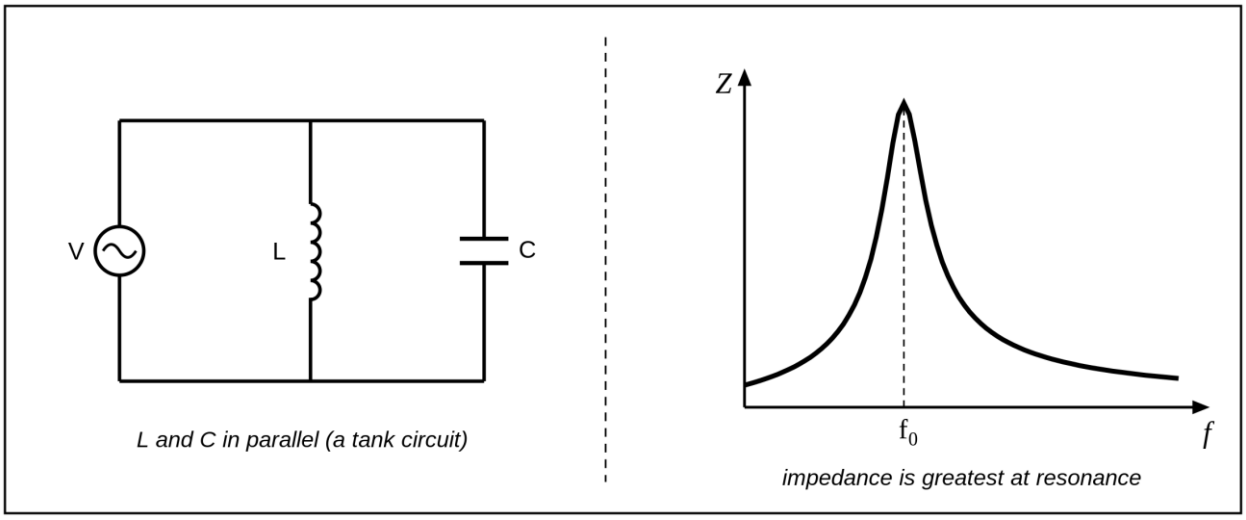


Figure: ***An inductor and a capacitor in parallel form a tank circuit.*** *At resonance the branch currents cancel, so the line current falls to a minimum while the impedance of the combination rises to a maximum, the exact opposite of the series case.*

The resonant frequency is given by the same formula, $f_0 = \frac{1}{2\pi\sqrt{LC}}$, but the behaviour is turned inside out. A series circuit at resonance has **minimum** impedance and **maximum** current; it eagerly accepts the resonant frequency and is called an **acceptor**. A parallel circuit at resonance has **maximum** impedance and **minimum** line current; it shuts that frequency out, and is called a **rejector**. Same two components, same resonant frequency, opposite job.

Resonance at a glance

Both circuits resonate at $f_0 = \frac{1}{2\pi\sqrt{LC}}$, where $X_L = X_C$.

Series (acceptor): impedance is minimum ($Z = R$), current is maximum ($I = \frac{V}{R}$), current in phase with voltage. Sharpness $Q = \frac{1}{R}\sqrt{\frac{L}{C}}$, with $V_L = V_C = QV$ and bandwidth $\Delta f = \frac{f_0}{Q}$.

Parallel (rejector): impedance is maximum, line current is minimum. The same f_0, the opposite behaviour.

Five facts, two circuits, and one capacitor about to hold a hundred times the voltage it was offered. Time to make resonance earn its keep. Four worked examples follow.

BINDER Example 28

A coil of inductance $200\mu H$ is connected in series with a capacitor of $200pF$. Find the resonant frequency of the combination.

Solution

The resonant frequency depends only on the product LC:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{200 \times 10^{-6}H \times 200 \times 10^{-12}F}} = 796kHz$$

Making Sense of the Answer: *Small inductances and capacitances give high resonant frequencies; radio work lives among microhenries and picofarads.*

Think Like a Physicist: *Only the product LC matters. To double the resonant frequency you must quarter that product, by shrinking either component.*

BINDER Example 29

A capacitor of $0.4\mu F$, a coil of inductance $0.4H$, a resistor of 10Ω and a lamp (whose resistance may be neglected) are joined in series with an alternating supply of $0.01V$ rms, whose frequency can be varied. Find (a) the resonant frequency, (b) the current at resonance, (c) the voltage across the capacitor at resonance, and (d) the Q-factor. Describe how the lamp's brightness changes as the frequency is swept through f_0.

Solution

(a) The resonant frequency is:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{0.4H \times 0.4 \times 10^{-6}F}} = 398Hz$$

(b) At resonance the reactances cancel, so the impedance is just the resistance, $Z = R = 10\Omega$, and the current is:

$$I = \frac{V}{R} = \frac{0.01V}{10\Omega} = 0.001A = 1mA$$

(c) The capacitor's reactance at resonance, using $\omega_0 = \frac{1}{\sqrt{LC}} = \frac{1}{\sqrt{0.4H \times 0.4 \times 10^{-6}F}} = 2500rad/s$, is:

$$X_C = \frac{1}{\omega_0 C} = \frac{1}{2500rad/s \times 0.4 \times 10^{-6}F} = 1000\Omega$$

So the voltage across it is:

$$V_C = IX_C = 1 \times 10^{-3}A \times 1000\Omega = 1.0V$$

(d) The Q-factor is:

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}} = \frac{1}{10\Omega}\sqrt{\frac{0.4H}{0.4 \times 10^{-6}F}} = 100$$

As the frequency rises from low to high, the lamp begins dim, brightens to a sharp maximum exactly at $398Hz$, then fades again. It is brightest at resonance, where the current is greatest.

Making Sense of the Answer: *The supply is a mere* $0.01V$*, yet* $1.0V$*, a full hundred times more, appears across the capacitor. That is voltage magnification by the factor* $Q = 100$*. The inductor carries the same* $1.0V$*, equal and opposite, so the two cancel and the supply sees only the* 10Ω *resistor.*

Think Like a Physicist: *At resonance reach straight for* $V_C = QV$*. Here* $Q = 100$ *and* $V = 0.01V$*, so* $V_C = 1.0V$ *in a single line, confirming the longer calculation.*

HOT Example 30

A tuned circuit resonates at $1.0MHz$ with a Q-factor of 125. (a) Find its bandwidth. (b) Two stations broadcast at $1.000MHz$ and $1.004MHz$. Can this circuit separate them?

Solution

(a) The bandwidth is the resonant frequency divided by the Q-factor:

$$\Delta f = \frac{f_0}{Q} = \frac{1.0 \times 10^6 Hz}{125} = 8kHz$$

(b) The two stations are separated by $1.004MHz - 1.000MHz = 4kHz$. Since this gap ($4kHz$) is smaller than the bandwidth ($8kHz$), both stations fall within the same resonance peak and cannot be cleanly separated. A higher Q, giving a narrower bandwidth, would be needed.

Making Sense of the Answer: *Bandwidth measures how fussy a circuit is. A bandwidth of* $8kHz$ *means the circuit passes a band* $8kHz$ *wide; stations closer together than that blur into one.*

Think Like a Physicist: *Selectivity and* Q *are the same idea seen from two sides. To pick apart close stations, raise* Q *and so narrow the bandwidth.*

HOT Example 31

A coil of inductance $2mH$ is connected in parallel with a capacitor of $8nF$ to form a tank circuit. (a) Find the resonant frequency. (b) State how the impedance and the supply current behave there, and name the circuit's role.

Solution

(a) The parallel circuit resonates at the same frequency as the series one:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{2 \times 10^{-3}H \times 8 \times 10^{-9}F}} = 39.8kHz$$

(b) At this frequency the currents in the two branches are equal and opposite, so the line current drawn from the supply falls to a minimum (ideally zero) while the impedance of the combination rises to a maximum. The circuit therefore shuts out the frequency f_0; it is acting as a rejector, or tank circuit.

Making Sense of the Answer: *The same* L *and* C*, merely rewired from series into parallel, resonate at the same frequency but do the opposite job: series accepts, parallel rejects.*

Think Like a Physicist: *The formula* $f_0 = \frac{1}{2\pi\sqrt{LC}}$ *is shared by both arrangements. What differs is whether the impedance bottoms out (series) or peaks (parallel) at resonance.*

And there it is: the circuit that listens for one frequency and shrugs at the rest. Kipute has gone home to announce that he can build a radio from a coil and a scrap of foil, which is very nearly true and entirely his teacher's fault. Kipanga, more careful, has noticed that the same L and C resonate at the same frequency whether wired in line or side by side, and that only the impedance cannot decide whether to vanish or to soar. He is exactly right.

THE TRANSFORMER

We opened this chapter with a quarrel: the war between one-way direct current and the back-and-forth of alternating current, and a promise that one quiet device would settle it. That device is the **transformer**, and its moment has arrived. A transformer takes an alternating voltage in at one value and hands it back out at another, larger or smaller, almost without loss and with no moving parts at all. It is the reason your phone charger can sip 5V from a 240V socket, and the reason the power born at Mtera can travel the length of the country without melting the wires on the way. It is also, as we shall see, the single reason the whole world chose alternating current.

How a transformer works: one flux, two coils

A transformer is two coils wound on the same iron core. One, the **primary**, is connected to the alternating supply; the other, the **secondary**, delivers the output. They are not joined by any wire; what links them is the iron.

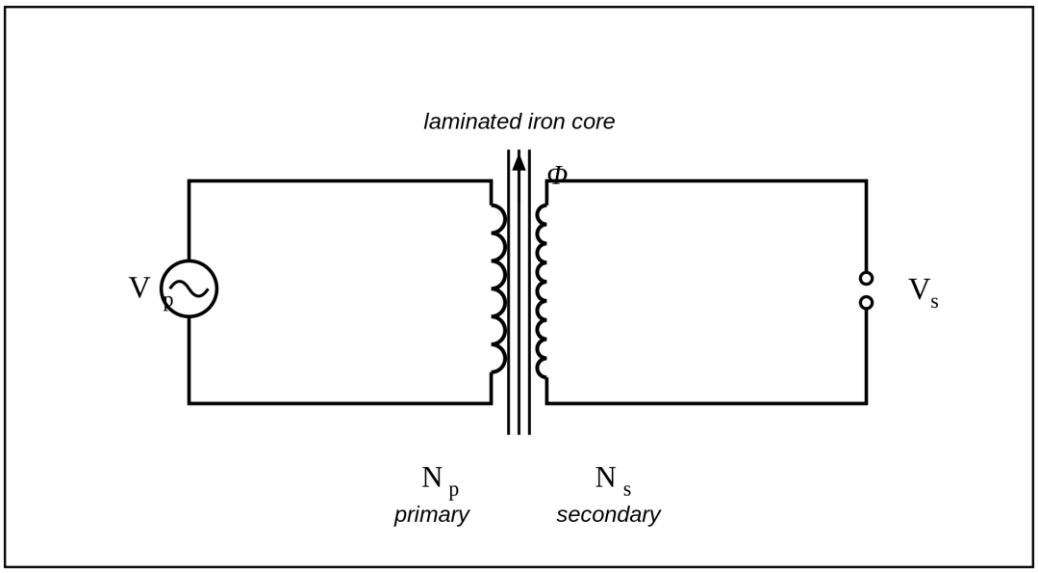


Figure: ***A transformer:*** *a primary coil and a secondary coil wound on a common laminated iron core. The alternating current in the primary sets up an alternating flux in the core, and that shared flux threads the secondary.*

Here is the chain of cause and effect. The alternating current in the primary sets up an alternating magnetic flux in the iron core. The core is a closed loop of iron, so it guides almost the whole of that flux around and through the secondary coil. Now recall the one fact we proved back in the alternator section: a coil of N turns threaded by a changing magnetic flux develops across itself a voltage equal to N times the rate at which the flux through one turn is changing. The very same changing flux Φ threads both coils, so each develops a voltage set by its own number of turns:

$$V_p = N_p \frac{d\Phi}{dt}, \qquad V_s = N_s \frac{d\Phi}{dt}$$

The rate of change of magnetic flux $\frac{d\Phi}{dt}$ is the same in both, so when we divide one equation by the other it cancels, leaving the **transformer equation**:

$$\frac{V_p}{V_s} = \frac{N_p}{N_s}$$

The voltages are in the same ratio as the turns. Give the secondary more turns than the primary ($N_s > N_p$) and it delivers a **higher** voltage; this is a **step-up** transformer. Give it fewer turns and the voltage comes out **lower**; this is a **step-down** transformer. The transformer's one trick is to trade turns for volts.

Current and power: what you gain in volts you pay in amps

A transformer cannot create energy. In an **ideal** transformer, one with no losses at all, the power fed into the primary equals the power drawn from the secondary:

$$V_p I_p = V_s I_s$$

Rearranging, and using the transformer equation, the currents come out in the **inverse** ratio of the turns:

$$\frac{I_p}{I_s} = \frac{V_s}{V_p} = \frac{N_s}{N_p}$$

This is the crucial trade. Step the voltage **up** and the current is stepped **down** in exactly the same proportion. A step-up transformer therefore delivers high voltage at low current, and that, as the last section of this topic will show, is precisely what long-distance transmission demands.

Notice one thing the transformer cannot do: work on steady direct current. A constant current makes a constant magnetic flux, and a magnetic flux that does not change induces no voltage in the secondary at all. Connect a battery to the primary and the secondary stays dead (save for a single flick at switch-on). A transformer lives entirely on the **changing** flux that only alternating current provides.

Why a real transformer falls short of perfection

A real transformer reaches only about 95% to 98% efficiency, never the full hundred. Its energy leaks away by three routes, and knowing them tells the engineer how to build a better one.

1) Copper loss

The windings are made of wire, and wire has resistance, so a current I through a winding of resistance R wastes I^2R as heat. Thicker, lower-resistance wire keeps this down.

2) Hysteresis loss

Every cycle, the alternating current drags the core's magnetisation right around a loop, the **B-H loop**, and forcing the iron round that loop costs energy. The energy lost per cycle, in each cubic metre of core, is exactly the **area enclosed by the loop**. The cure is to build the core from a soft magnetic material whose loop is as narrow as possible.

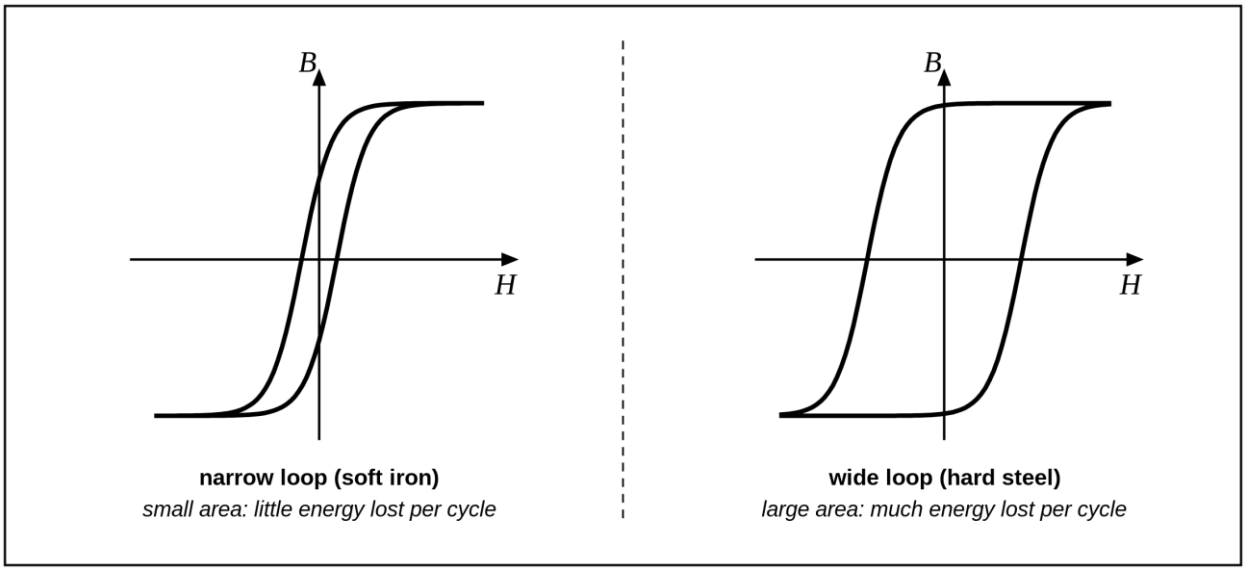


Figure: ***The B-H loop of the core material.*** *The energy turned to heat each cycle equals the area inside the loop. Soft iron (left) has a narrow loop and wastes little; hard steel (right) has a fat loop and wastes a great deal. A transformer core is therefore made of soft iron.*

This is why an examiner who asks ***"why should the core have a narrow hysteresis loop?"*** is really asking you to say: *because the loop area is the energy lost per cycle, and a narrow loop means a small area, so little heat and a more efficient transformer.*

3) Eddy-current loss

The changing flux does not only thread the coils; it also threads the iron of the core itself, and so drives little circulating currents, called **eddy currents**, round and round inside the metal. These currents heat the core for no useful purpose. The cure is to build the core not as one solid block but from many thin sheets, **laminations**, each lightly insulated from its neighbours, so the circulating currents are chopped into small, feeble loops.

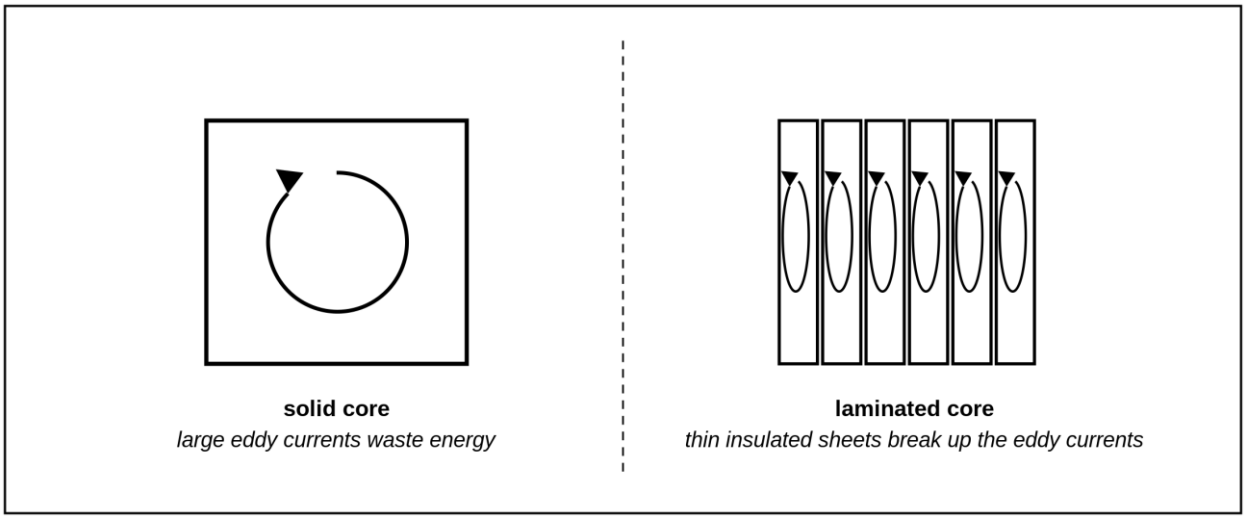


Figure: ***A solid core (left) lets one large eddy current circulate and waste energy.*** *Building the core from thin insulated laminations (right) breaks the flow into many small, weak loops, cutting the loss sharply.*

So the ideal core is soft iron of high permeability (to guide the flux well), with a narrow loop (little hysteresis) and built of thin laminations (little eddy-current loss). With these three cares taken, the small leftover gives the 95% to 98% efficiency we quoted. Efficiency itself is just:

$$\eta = \frac{P_{out}}{P_{in}} \times 100\%$$

Why the grid is alternating: from Mtera to your socket

Now the promise of the opening pays off. The alternators at the Mtera hydroelectric plant generate electricity at a few thousand volts. To carry it hundreds of kilometres, the heating loss I^2R in the lines must be kept tiny, and since the cables' resistance R is fixed, the only way is to make the current I small. By the transformer's trade, small current means high voltage. So a step-up transformer at the station lifts the voltage to the transmission level (the TANESCO grid runs at 132kV and above), the lines then carry the power at high voltage and low current with very little loss, and step-down transformers near the towns, such as at a Tanga substation, bring it back down, stage by stage, to the 240V at your wall socket.

Every link in that chain is a transformer, and a transformer works only on alternating current. That, in one sentence, is why the world's grids are a.c. and why the device cast as the hero of our opening truly earned the part.

Four worked examples to put the hero through its paces. They run from a one-line turns ratio up to the transmission sum that explains the whole grid, with the national-examination mutual-inductance item sitting in the middle. Kipute, who once tried to charge his phone straight from the wall and has the scorched memory to prove it, is unusually attentive.

BINDER Example 32

A transformer steps the 240V mains down to 12V for a doorbell. If the primary has 2000 turns, how many turns has the secondary?

Solution

The turns follow the voltages, so:

$$N_s = N_p\left(\frac{V_s}{V_p}\right) = 2000 \times \frac{12V}{240V} = 100 \text{ turns}$$

Making Sense of the Answer: *The voltage was cut to a twentieth* (12V *from* 240V)*, so the turns are cut to a twentieth as well, from* 2000 *to* 100*. Fewer turns on the secondary, lower voltage out: a step-down transformer.*

Think Like a Physicist: *The turns ratio and the voltage ratio are one and the same. Find one and you have the other, with no physics left to do.*

HOT Example 33

A step-down transformer runs a 12V, 24W lamp from the 240V mains. Assuming it is ideal, find (a) the current in the lamp (the secondary current), and (b) the current drawn from the mains (the primary current).

Solution

(a) The lamp's current follows from its power rating, $P = V_sI_s$:

$$I_s = \frac{P}{V_s} = \frac{24W}{12V} = 2A$$

(b) An ideal transformer conserves power, $V_pI_p = V_sI_s$, so:

$$I_p = \frac{V_sI_s}{V_p} = \frac{12V \times 2A}{240V} = 0.1A$$

Making Sense of the Answer: *The secondary delivers $2A$ at $12V$, the primary takes only $0.1A$ at $240V$. Both deliver the same $24W$. The transformer stepped the voltage down by twenty and stepped the current up by twenty, exactly preserving the power.*

Think Like a Physicist: *Whenever a transformer is called ideal, anchor on $V_pI_p = V_sI_s$. Voltage and current always move in opposite directions, so the watts stay put.*

REAL Example 34

A power station delivers $11MW$ along a line of total resistance 5Ω. Compare the power wasted as heat in the line when the power is sent at $11kV$ with the loss when a transformer first steps it up to $132kV$.

Solution

The line current is $I = \frac{P}{V}$, and the heating loss is I^2R. At $11kV$:

$$I = \frac{11 \times 10^6W}{11 \times 10^3V} = 1000A, \quad I^2R = (1000A)^2 \times 5\Omega = 5MW$$

At $132kV$, twelve times the voltage, the current is twelve times smaller:

$$I = \frac{11 \times 10^6W}{132 \times 10^3V} = 83.3A, \quad I^2R = (83.3A)^2 \times 5\Omega = 35kW$$

The loss falls from $5MW$ (almost half the power thrown away as heat) to a mere $35kW$, well under one per cent.

Making Sense of the Answer: *Raising the voltage twelve times cut the current twelve times, and the loss, which goes as I^2, fell by $12^2 = 144$ times. This single calculation is the whole case for high-voltage transmission, and so for the transformer, and so for a.c.*

Think Like a Physicist: *To move power a long way, push the voltage up and the current down. The loss I^2R rewards every extra volt handsomely, because it depends on the square of the current.*

HOT Example 35

A transformer draws $500W$ from the mains. Its copper loss is $12W$ and its iron loss (hysteresis and eddy currents together) is $8W$. Find (a) the output power and (b) the efficiency.

Solution

(a) The output is the input less the losses:

$$P_{out} = P_{in} - \text{losses} = 500W - 12W - 8W = 480W$$

(b) The efficiency is:

$$\eta = \frac{P_{out}}{P_{in}} \times 100\% = \frac{480W}{500W} \times 100\% = 96\%$$

Making Sense of the Answer: *A 96% efficiency sits right in the usual range. The $12W$ copper loss would shrink with thicker wire; the $8W$ iron loss is held down by the narrow-loop core and the laminations we met above.*

Think Like a Physicist: *Efficiency is always useful output over total input. For a transformer the only thing standing between the two is the sum of copper and iron losses; account for those and the rest comes straight through.*

From a coil spinning at Mtera to the socket on your wall, every step of the journey is a transformer quietly trading volts for amps. Kipute, having grasped at last why one cannot charge a phone straight from the wall, has declared the transformer "the most polite machine in physics, because it never asks for more power than it gives." Kipanga, checking the sums, confirms it gives back all it takes save a few honest watts of warmth, and even those, he notes, are kept small by a narrow loop and a well-laminated core.

POWER IN AC CIRCUITS AND POWER FACTOR

For a plain resistor, power is the simplest thing in the world: $P = I^2R$, and there is an end of it. But the moment an inductor or a capacitor enters the circuit, the voltage and the current fall out of step, and the word "power" turns slippery. We glimpsed the strangeness back in the pure-circuits section, where a perfect inductor or capacitor was found to dissipate, on average, **no power at all**. Now we make that precise, and then do something the examiners love and the accountants love even more: we turn it into money, and explain why a workshop in Dar es Salaam can be charged more for running the very same machines.

Average power: only the in-step part counts

At every instant the power delivered is the product of the instantaneous voltage and current, $p = vi$. When the two are in step, as in a resistor, this product is always positive: energy flows steadily into the component. But when they are a quarter cycle ($\pi/2$) apart, as in a pure reactance, the product is positive for half of each cycle and negative for the other half. Energy pours

in, then pours straight back out, and over a whole cycle the **average** comes to exactly zero. That is the precise reason a pure inductor or capacitor consumes no power: it borrows energy and returns it, every cycle, keeping nothing.

Take the voltage and current with a phase difference φ between them, $v = V_0 \sin\omega t$ and $i = I_0 \sin(\omega t - \varphi)$. The instantaneous power is their product:

$$p = vi = V_0 I_0 \sin\omega t \sin(\omega t - \varphi)$$

The product of two sines can be reshaped by the standard identity:

$$\sin A \sin B = \frac{1}{2}[\cos(A - B) - \cos(A + B)],$$

which turns it into:

$$p = \frac{1}{2} V_0 I_0 [\cos\varphi - \cos(2\omega t - \varphi)]$$

The first term inside the bracket is a constant; the second oscillates at twice the supply frequency, so over one complete cycle it averages to zero. The average power is therefore the constant part alone:

$$P = \frac{1}{2} V_0 I_0 \cos\varphi$$

Finally, since $V_{rms} = \frac{V_0}{\sqrt{2}}$ and $I_{rms} = \frac{I_0}{\sqrt{2}}$:

$$V_0 = V_{rms}\sqrt{2} \text{ and } I_0 = I_{rms}\sqrt{2}$$

Substituting to $P = \frac{1}{2} V_0 I_0 \cos\varphi$:

$$P = \frac{1}{2} \times V_{rms}\sqrt{2} \times I_{rms}\sqrt{2} \times \cos\varphi = \frac{1}{2} \times 2 \times V_{rms} \times I_{rms} \times \cos\varphi = V_{rms} I_{rms} \cos\varphi$$

The result is **average power** (also called the **true power)**, measured in watts:

$$\mathbf{P = V_{rms} I_{rms} \cos\varphi}$$

There is the origin of the $\cos\varphi$: it is the cosine of the phase angle between voltage and current, and it drops out of the algebra the moment the two are allowed to fall out of step. For a resistor $\varphi = 0$ and $\cos\varphi = 1$, so all of the product does real work; for a pure inductor or capacitor $\varphi = 90°$ and $\cos\varphi = 0$, so none of it does, exactly confirming the zero-power result we began with. (The $\cos(2\omega t - \varphi)$ term we threw away is the energy sloshing in and out at twice the supply frequency, the very oscillation an inductor and capacitor live by.)

The power triangle: true, reactive, and apparent power

The plain product of the rms voltage and current is itself given a name. It is the **apparent power**, written S and measured in **volt-amperes** (VA), because it is what the supply must appear to provide:

$$S = V_{rms} I_{rms}$$

Of this, the part actually turned into work or heat is the **true power** $P = S\cos\varphi$, in watts (W). The remaining part, $Q = S\sin\varphi$, merely surges in and out of the inductors and capacitors and does no net work; it is the **reactive power**, measured in **volt-amperes reactive (VAR).** Because true and reactive power are at right angles, like the resistance and reactance that cause them, the three combine into a right-angled **power triangle**.

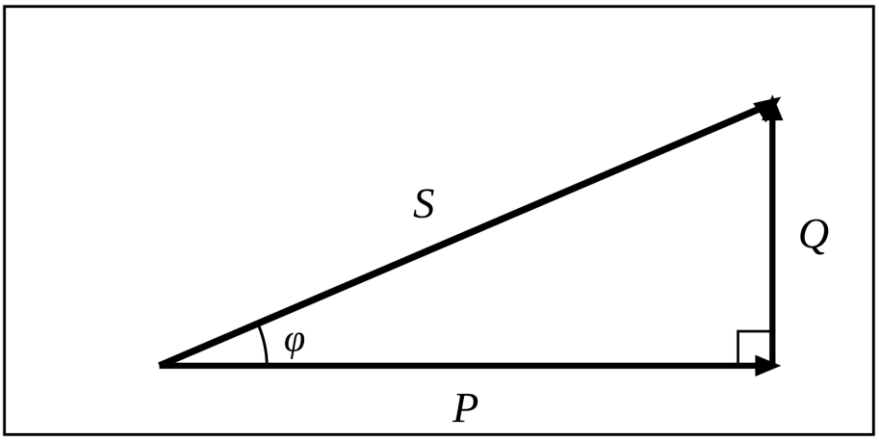


Figure: ***The power triangle.*** *The apparent power S is the hypotenuse, the true power P the base, and the reactive power Q the upright side. The angle between S and P is the same phase angle φ as in the impedance triangle.*

With S the hypotenuse:

$$S^2 = P^2 + Q^2$$

That is:

$$(V_{rms} I_{rms})^2 = (V_{rms} I_{rms} \cos\varphi)^2 + (V_{rms} I_{rms} \sin\varphi)^2$$

Power factor: how much of the apparent power works

The ratio of the true power to the apparent power is the **power factor**:

$$\textbf{power factor} = \cos\varphi = \frac{\mathbf{P}}{\mathbf{S}} = \frac{\mathbf{R}}{\mathbf{Z}}$$

The last equality follows from the impedance triangle of the series section: there the resistance R is the side adjacent to the angle φ and the impedance Z is the hypotenuse, so the cosine of φ, adjacent over hypotenuse, is exactly $\frac{R}{Z}$.

A power factor of 1 means a purely resistive circuit, with every volt-ampere doing real work; a power factor of 0 means a purely reactive one, with none of it doing any. When the current **lags** the voltage, as in an inductive circuit, the power factor is called **lagging**; when it **leads**, as in a capacitive one, it is called **leading**.

Why a poor power factor costs money

Most industrial loads, induction motors, transformers, the chokes in fluorescent lights, are strongly inductive, and run at a power factor of only about 0.7 to 0.8 lagging. Here is the trouble. To deliver a true power P at voltage V_{rms}, the supply must push a current,

$$I_{rms} = \frac{P}{V_{rms}\cos\varphi}$$

So the smaller the power factor, the **larger** the current needed for the very same useful power. That bloated current heats the supply cables more (I^2R again), and forces the utility to install thicker cables and bigger transformers to carry it. For this reason TANESCO bills large consumers on the **apparent** power in kVA, not merely the true power in kW. A workshop with a poor power factor draws extra current it never turns into work, and pays for the privilege.

Power-factor correction: the parallel capacitor

The cure is cheap and elegant. Connect a capacitor, or a bank of them, in **parallel** with the inductive load. The capacitor draws a leading current that cancels part of the inductor's lagging reactive current, so the net reactive power drops, the power factor climbs back toward 1, and the line current shrinks, all while the load goes on doing exactly the same real work.

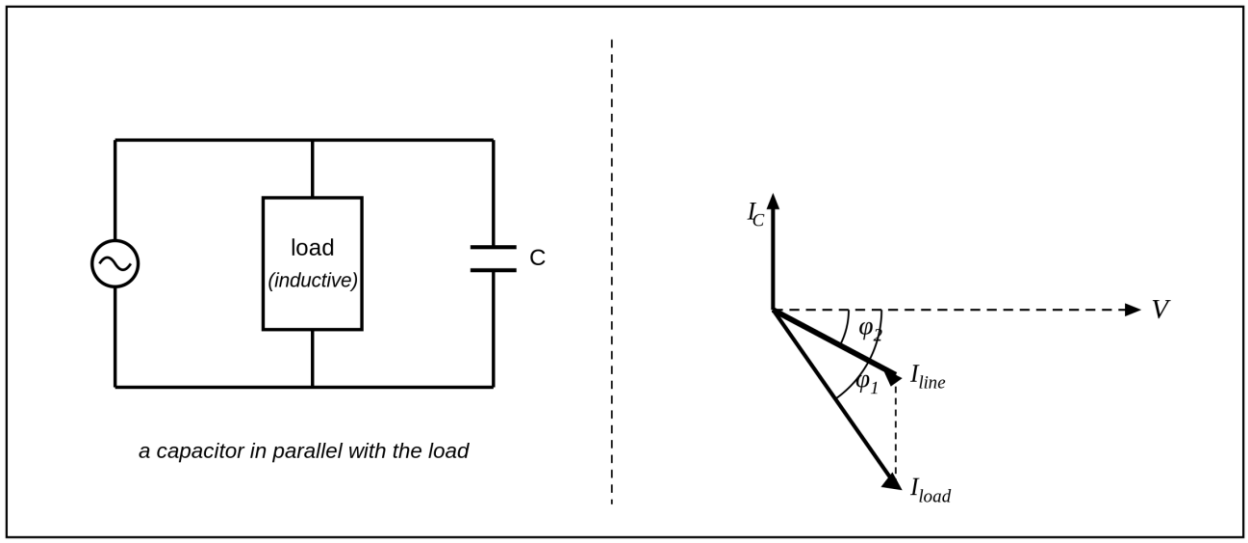


Figure: *Left: a capacitor placed in parallel with an inductive load. Right: the load current lags the voltage by a large angle; adding the capacitor's leading current pulls the line current up to a smaller angle and a smaller size, raising the power factor.*

Here is the calculation behind it. From the power triangle the reactive and true powers are linked by:

$$\tan\varphi = \frac{Q}{P}$$

So a load taking true power P at phase angle φ also draws reactive power $Q = P\tan\varphi$.

To lift the factor from $\cos\varphi_1$ to a better $\cos\varphi_2$ while holding the true power P fixed, the reactive power demanded from the supply must fall from $P\tan\varphi_1$ to $P\tan\varphi_2$, and the capacitor must make up the difference:

$$Q_C = P\tan\varphi_1 - P\tan\varphi_2 = P(\tan\varphi_1 - \tan\varphi_2)$$

A capacitor of its own draws reactive power:

$$Q_C = \frac{V_{rms}^2}{X_C} = \frac{V_{rms}^2}{\frac{1}{2\pi fC}} = V_{rms}^2 \times 2\pi fC$$

Solving for C:

$$\mathbf{C = \frac{Q_C}{2\pi f V_{rms}^2} = \frac{P(\tan\varphi_1 - \tan\varphi_2)}{2\pi f V_{rms}^2}}$$

The motor still turns, the lights still glow; only the wasteful reactive current that used to slosh all the way back to the power station is now supplied locally, by the capacitor, a few centimetres away. (The deeper reason the grid carries all this on alternating current remains the transformer of the previous section; power-factor correction is the same fight against wasted current, waged at the factory instead of the substation.)

Four worked examples, running from a one-line power factor up to a capacitor sized to save a workshop money. Kipanga, who has just learned that his uncle's welding shop is billed in kVA, is taking notes with unusual seriousness.

BINDER Example 36

A coil has resistance 12Ω and, at the working frequency, an impedance of 20Ω. Find its power factor and state whether it is leading or lagging.

Solution

The power factor is the ratio of resistance to impedance:

$$\cos\varphi = \frac{R}{Z} = \frac{12\Omega}{20\Omega} = 0.60$$

A coil is inductive, so the current lags the voltage: the power factor is 0.60 **lagging**.

Making Sense of the Answer: *Only* 60% *of the apparent power does real work. The other* 40% *is reactive, sloshing in and out of the coil's magnetic field each cycle. A bare coil is already a fairly poor power factor.*

Think Like a Physicist: *Power factor is just* $\frac{R}{Z}$*, the cosine of the impedance triangle's angle. Resistive part over total opposition; nothing more to memorise.*

REAL Example 37

Two workshops each draw the same useful power of 10kW from the 415V mains. The first has a power factor of 1.0; the second, full of motors, runs at 0.70 lagging. Explain, with figures, why TANESCO charges the second workshop more even though both "use the same machines."

Solution

The current each must draw is $I_{rms} = \frac{P}{V_{rms}\cos\varphi}$.

For the first workshop:

$$I_{rms} = \frac{10000\text{W}}{415\text{V} \times 1.0} = 24.1\text{A}$$

For the second workshop:

$$I_{rms} = \frac{10000\text{W}}{415\text{V} \times 0.70} = 34.4\text{A}$$

Both turn the same 10kW into useful work, but the second draws 34.4A against the first's 24.1A, about 43% more current, for nothing extra in return. That surplus current heats TANESCO's cables and ties up its transformers, so TANESCO meters the **apparent** power: the first workshop presents 10kVA, the second 14.3kVA $\left(\text{from } \frac{10\text{kW}}{0.7}\right)$. The second pays for capacity it wastes.

Making Sense of the Answer: *"The same machines" do the same real work, true, but the poor power factor makes the second workshop a heavier customer on the network. The bill follows the current, and the current follows the power factor.*

Think Like a Physicist: *Whenever someone says two loads are "the same," check the current each actually draws. Real power can be equal while apparent power, and therefore cost, differs widely.*

HOT Example 38

A 50Ω resistor, a 0.100H inductor and a $10.0\mu\text{F}$ capacitor are connected in series to a 60Hz source. The rms current in the circuit is 2.75A. Find (a) the rms voltage of the source, (b) the power factor, and (c) the average power.

Solution

(a) First the reactances at $\omega = 2\pi f = 2\pi \times 60\text{Hz} = 377\text{rad/s}$:

$$X_L = \omega L = 377\text{rad/s} \times 0.100\text{H} = 37.7\Omega, \quad X_C = \frac{1}{\omega C} = \frac{1}{377\text{rad/s} \times 10.0 \times 10^{-6}\text{F}} = 265\Omega$$

The impedance, with the reactances opposing, is:

$$Z = \sqrt{R^2 + (X_L - X_C)^2} = \sqrt{(50\Omega)^2 + (37.7\Omega - 265\Omega)^2} = 233\Omega$$

So the source voltage is:

$$V_{rms} = I_{rms}Z = 2.75\text{A} \times 233\Omega = 641\text{V}$$

(b) The power factor is

$$\cos\varphi = \frac{R}{Z} = \frac{50\Omega}{233\Omega} = 0.215 \text{ (leading)}$$

Leading, because the capacitor's reactance wins and the circuit is net capacitive.
(c) The average power is dissipated only in the resistor:

$$P = I_{rms}^2 R = (2.75A)^2 \times 50\Omega = 378W$$

Making Sense of the Answer: *The circuit is far off resonance, so its impedance* (233Ω) *is much larger than the resistance* (50Ω) *and the power factor is poor. Of the* $641V \times 2.75A = 1763VA$ *of apparent power, only* $378W$ *does real work, just as* 0.215 *of it should.*

Think Like a Physicist: *Average power is always* $I_{rms}^2 R$*, the resistor's alone, no matter how much the reactances inflate the voltage. Find Z for the current, but find R for the power.*

HOT Example 39

A Dar es Salaam workshop draws 50kW at a power factor of 0.70 lagging from the 415V, 50Hz mains. What capacitance, placed in parallel, will raise the power factor to 0.95?

Solution

The capacitor must supply the reactive power $Q_C = P(\tan\varphi_1 - \tan\varphi_2)$.
From $\cos\varphi_1 = 0.70$ we get $\tan\varphi_1 = 1.02$, and from $\cos\varphi_2 = 0.95$, $\tan\varphi_2 = 0.329$. So:

$$Q_C = 50\text{kW} \times (1.02 - 0.329) = 34.6\text{kVAR}$$

And the capacitance that provides this at 415V, 50Hz is:

$$C = \frac{Q_C}{2\pi f V_{rms}^2} = \frac{34600\text{VAR}}{2\pi \times 50\text{Hz} \times (415\text{V})^2} = 6.39 \times 10^{-4}\text{F} = 639\mu\text{F}$$

As a check, the line current falls from $I = \frac{50000\text{W}}{415\text{V}\times 0.70} = 172\text{A}$ before correction to $\frac{50000\text{W}}{415\text{V}\times 0.95} = 127\text{A}$ after, a saving of 45A for the same 50kW of work.

Making Sense of the Answer: *A* $639\mu F$ *bank trims the current by a quarter without touching the workshop's output. The capacitor supplies the motors' reactive power on the spot, so TANESCO's cables no longer have to carry it back and forth.*

Think Like a Physicist: *Correction is bookkeeping in reactive power: find* Q *before and after from* $P\tan\varphi$*, give the difference to a capacitor, and size it from* $Q_C = V_{rms}{}^2 \times 2\pi fC$*. The real power never enters the capacitor's job.*

Power, it turns out, keeps two sets of books: the watts that do the work and the volt-amperes the supply must carry to deliver them. Keep the two close together, a power factor near one, and the network is content and the bill is small. Let them drift apart and the wires run hot for nothing. Kipute has resolved to fit his uncle's welding shop with "a bucket of capacitors," and Kipanga, for once, cannot fault the idea, only recommend that he first work out the right number of microfarads.

RECTIFICATION AND THE MEASUREMENT OF AC

Alternating current is unbeatable for generating and transmitting power, as the transformer made plain. Yet almost every piece of electronics that the power finally reaches, a phone, a laptop, the radio, the CT scanner at Bugando Hospital, runs on steady **direct** current inside. Somewhere between the wall socket and the chip, the back-and-forth of a.c. must be turned into the one-way flow of d.c. The device that does it is the **rectifier**, and it is built from diodes. We close the chapter's machinery with two questions: *how do we convert a.c. into d.c., and how, since an ordinary meter is fooled by a current that keeps reversing, do we measure a.c. in the first place?*

Everything rests on one fact about the **diode**, which you met among the semiconductors: a diode is a one-way valve for current. It conducts freely when the voltage pushes it the "forward" way, and blocks almost completely when the voltage tries to push it backward. Feed alternating current into a diode and only the forward halves get through.

The half-wave rectifier

The simplest rectifier is a single diode in series with the load. During the half-cycle that drives the diode forward, current flows and the load receives the voltage; during the reverse half-cycle the diode blocks, no current flows, and the output is zero. The result is a string of positive humps with empty gaps between them.

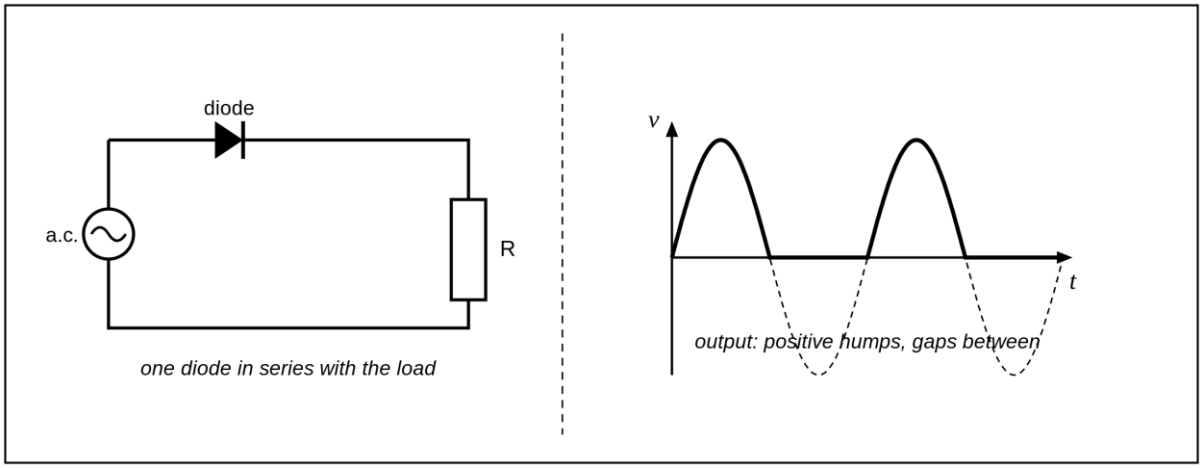


Figure: ***A half-wave rectifier:*** *a single diode in series with the load. Only the half-cycles that bias the diode forward reach the load, so the output is a row of positive humps separated by gaps where the diode blocks.*

What steady (d.c.) value does this pulsing output average to? We must average the output over a whole period. For the first half, from $\theta = 0$ to π, the output is $V_0 \sin\theta$; for the second half it is zero. So:

$$V_{mean} = \frac{1}{2\pi}\int_0^{\pi} V_0 \sin\theta \, d\theta = \frac{V_0}{2\pi}[-\cos\theta]_0^{\pi} = \frac{V_0}{2\pi} \times (1+1) = \frac{V_0}{\pi} = 0.318V_0$$

Therefore:

$$\mathbf{V_{mean} = \frac{V_0}{\pi} = 0.318V_0}$$

This is the average value of a half-wave rectified sine wave over a complete cycle. Half-wave rectification is simple but wasteful: it throws away every other half-cycle and leaves a very bumpy output.

The full-wave bridge rectifier

To use **both** halves of the wave, we arrange four diodes in a **bridge**. On either half-cycle, two of the four diodes conduct and steer the current through the load in the **same** direction, while the other two block. The negative halves are not discarded but flipped up to join the positive ones, so the output is an unbroken row of positive humps with no gaps.

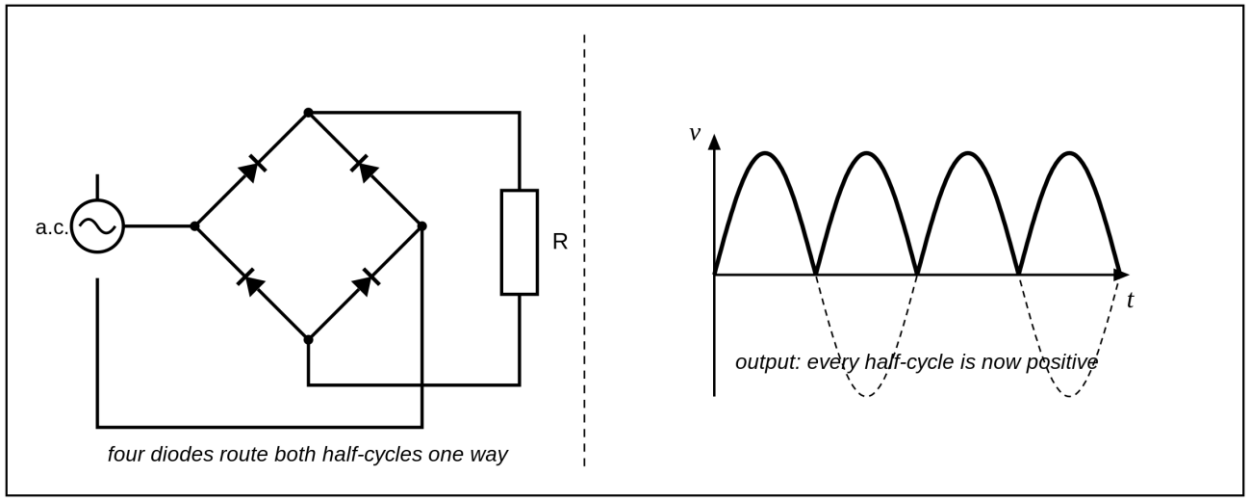


Figure: ***A full-wave bridge rectifier:*** *four diodes route both half-cycles through the load in one direction. Whichever way the supply pushes, the current crosses the load the same way, so every half-cycle appears as a positive hump.*

Now every half-cycle contributes, and the pattern repeats every π, so we average over just half a period:

$$V_{mean} = \frac{1}{\pi}\int_0^{\pi} V_0 \sin\theta \, d\theta = \frac{V_0}{\pi}[-\cos\theta]_0^{\pi} = \frac{V_0}{\pi} \times 2 = \frac{2V_0}{\pi} = 0.637V_0$$

Therefore:

$$\mathbf{V_{mean} = \frac{2V_0}{\pi} = 0.637V_0,}$$

exactly twice the half-wave value. The full-wave output is smoother and wastes nothing, which is why it is the arrangement found inside almost every real power supply.

Smoothing: the reservoir capacitor

Even the full-wave output still rises and falls between zero and the peak. To smooth it into something close to steady d.c., a large capacitor is connected in **parallel** with the load. It charges up to the peak of each hump, then, during the dips, discharges slowly through the load, holding the voltage up and filling in the troughs. What is left is a nearly steady voltage with only a small **ripple** riding on top.

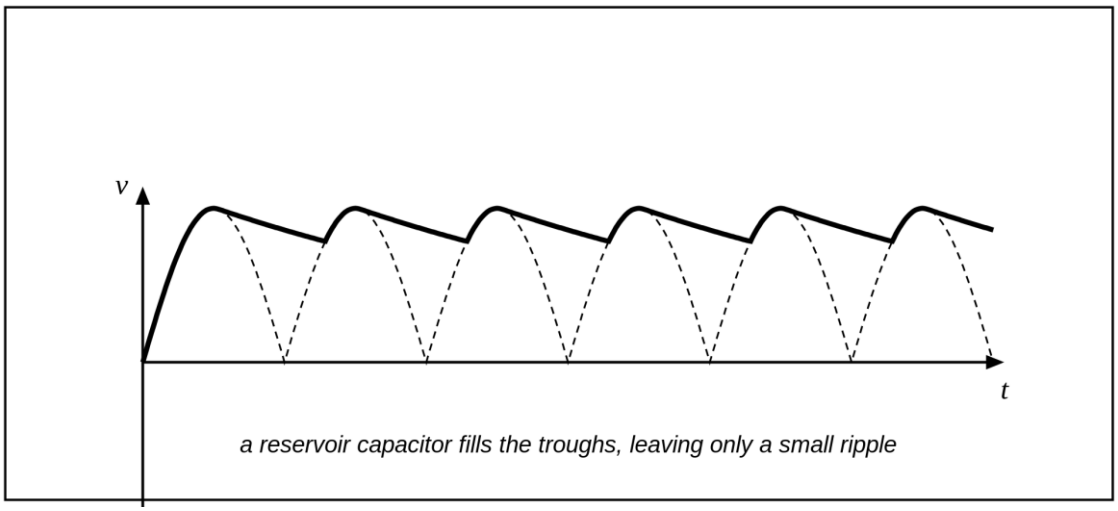


Figure: *With a reservoir capacitor across the load, the output charges to each peak and discharges only a little before the next hump arrives. The deep troughs are filled in, leaving a near-steady d.c. voltage with a small ripple.*

The larger the capacitor (and the larger the load resistance), the slower the discharge and the smaller the ripple. This is precisely how a phone charger or the d.c. supply inside the Bugando scanner turns the 240V a.c. mains, once a transformer has stepped it down, into the smooth d.c. its electronics demand.

Measuring alternating current

A **moving-coil meter**, the workhorse of the d.c. laboratory, measures the **average** current through it. On alternating current the average over a full cycle is zero, so the needle does not move: such a meter simply cannot read a.c. To measure a.c. we need an instrument that responds to the current's **size** regardless of its direction. Three designs do this.

1) **The rectifier-type meter** rectifies the current first, then lets an ordinary moving-coil meter read the mean of the rectified wave. Since that mean is not the rms value, the scale is marked up by a fixed **form factor**. For a sine wave the form factor is the ratio of the rms value to the full-wave mean, $\frac{\text{I}_{\text{rms}}}{\text{I}_{\text{dc}}} = \frac{\pi}{2\sqrt{2}} = 1.11$, so the scale reads 1.11 times the mean and thereby shows rms directly.
2) **The moving-iron meter** works by the force between pieces of iron magnetised by the current; that force depends on the **square** of the current, so the instrument responds to the average of I^2 and reads the rms value straight off, on a.c. or d.c. alike.
3) **The hot-wire meter** uses the heating of a fine wire, I^2R, which likewise depends on the square of the current; it too reads true rms, whatever the shape of the wave. Both (2 and 3 designs) have crowded, non-uniform scales, the price of squaring.

Three worked examples to put the diodes to work.

BINDER Example 40

A half-wave rectifier supplies 50V d.c. to a resistive load of 800Ω. The diode has a resistance of 25Ω. Calculate the peak a.c. voltage required.

Solution

The d.c. delivered to the load fixes the d.c. (mean) current through it:

$$\text{I}_{\text{dc}} = \frac{\text{V}_{\text{dc}}}{\text{R}} = \frac{50\text{V}}{800\Omega} = 0.0625\text{A}$$

For a half-wave rectifier the mean current is $1/\pi$ of the peak, $\text{I}_{\text{dc}} = \frac{\text{I}_0}{\pi}$, so the peak current is:

$$\text{I}_0 = \pi\text{I}_{\text{dc}} = \pi \times 0.0625\text{A} = 0.196\text{A}$$

This peak current is driven by the peak voltage through the diode and the load in series:

$$\text{R} + \text{r} = 800\Omega + 25\Omega = 825\Omega$$

It follows that:

$$\text{V}_0 = \text{I}_0(\text{R} + \text{r}) = 0.196\text{A} \times 825\Omega = 162\text{V}$$

Making Sense of the Answer: *A hefty* $162V$ *peak is needed to leave* $50V$ *of steady d.c., because half-wave rectification keeps only a thin slice of each cycle, and the mean of that slice is just* 0.318 *of the peak.*

Think Like a Physicist: *Work inward from the load. The d.c. it needs sets* I_{dc}*; the factor* π *lifts that to the peak* I_0*; and Ohm's law across the whole series path then gives the peak voltage.*

REAL Example 41

Inside a charger, a transformer steps the $240V$ mains down to a peak of $12V$, and a full-wave bridge then rectifies it. (a) Find the d.c. (mean) output voltage before smoothing. (b) State what happens when a large reservoir capacitor is added across the load.

Solution

(a) For a full-wave rectifier the mean is 0.637 of the peak:

$$V_{mean} = \frac{2V_0}{\pi} = \frac{2 \times 12V}{\pi} = 7.64V$$

(b) The capacitor charges up to the peak of each hump and discharges only slightly before the next arrives. The output therefore rises from this $7.64V$ average towards the peak $12V$, holding nearly steady at that level with only a small ripple.

Making Sense of the Answer: *Without smoothing the output averages* $7.64V$ *but swings all the way down to zero between humps; with a good capacitor it sits close to the* $12V$ *peak, which is what the charger's electronics actually want.*

Think Like a Physicist: *Rectification decides the* ***shape****; the capacitor decides the* ***smoothness****. The mean* $0.637V_0$ *is the starting point, and smoothing pushes the working voltage up toward the peak.*

HOT Example 42

An alternating current of peak value $2.0A$ flows in turn through three meters: a moving-coil meter, a rectifier-type meter, and a hot-wire meter. What does each one read?

Solution

The **moving-coil** meter reads the average current over a full cycle. For a symmetric a.c. that average is zero, so it reads:

$$\text{moving} - \text{coil reading} = 0$$

The **rectifier** meter reads the full-wave mean, $I_{dc} = \frac{2I_0}{\pi} = 0.637 \times 2.0A = 1.27A$, and its scale multiplies this by the form factor 1.11 to display the rms value:

$$\text{rectifier reading} = 1.11 \times 1.27A = 1.41A$$

The **hot-wire** meter responds to I^2 and so reads the true rms value directly:

$$I_{rms} = \frac{I_0}{\sqrt{2}} = \frac{2.0A}{\sqrt{2}} = 1.41A$$

So the readings are 0, $1.41A$ and $1.41A$ respectively.

Making Sense of the Answer: *The rectifier meter and the hot-wire meter agree, both showing the rms value* $1.41A$*, while the moving-coil meter, blind to a reversing current, reads nothing.*

Think Like a Physicist: *Ask what each meter physically responds to. Average current, and a.c. averages to zero; the square of the current, and you get rms. The rectifier meter is the in-between case, rescued by the* 1.11 *form factor.*

With that, the toolkit is complete. We can make alternating current, measure it, picture it with phasors, send it through resistors, inductors, and capacitors alone and combined, tune it to resonance, raise and lower it with transformers, reckon its power, and finally tame it into the steady d.c. that electronics drink. All that remains is to watch the whole apparatus go to work in the world, which is exactly where the chapter turns next.

APPLICATIONS OF AC THEORY

Alternating current is at its most convincing when it stops being equations and becomes objects. This closing section takes eight devices of everyday Tanzanian life and reads each one through the physics the chapter has built, from the national grid that spans the country down to the charger on the bedside table. Nothing new is introduced here; the aim is only to watch the theory already in hand earn its keep in the world.

1. The national grid

The whole chapter, in a sense, exists for this. The alternators at the Mtera and Kidatu power stations turn mechanical energy into an alternating EMF; to carry that power hundreds of kilometres without wasting it as heat in the lines, the current must be small, and so, by the transformer's trade of $\frac{V_p}{V_s}$ against turns, the voltage must be very large. A step-up transformer at the station lifts it to the transmission voltage; the line carries the power at high voltage and low current with little I^2R loss; and step-down transformers near the towns bring it back to a voltage safe for the home. The entire chain works only because the transformer works, and the transformer works only on alternating current.

The TANESCO grid transmits at 132kV, and on the newer southern interconnector at 220kV, stepping down at substations such as Tanga's to 33kV and 11kV and at last to the 240V of the wall socket. As the transmission example of the power section showed, sending the same power at 11kV instead of 132kV would multiply the line loss by 144 and melt the economics of the grid. High-voltage a.c. transmission is not a convenience but the only way bulk power can travel across a country the size of Tanzania.

2. Tuning a radio or television

Inside every receiver sits a resonant circuit, an inductor and a variable capacitor, that responds most strongly at the one frequency $f_0 = \frac{1}{2\pi\sqrt{LC}}$. The aerial picks up every station at once, but only the station whose frequency matches f_0 drives the circuit at resonance and is magnified; the rest, off resonance, are passed over. Turning the tuning knob rotates the plates of the variable capacitor, changing C and sliding f_0 from one station to the next. The sharper the resonance, the higher the circuit's Q, the more cleanly one station is separated from its neighbour.

A typical receiver pairs a coil of about 250µH with a capacitor variable from 20pF to 400pF, giving a tuning range that comfortably spans the 540kHz to 1610kHz of the AM band on which Radio Free Africa and TBC reach their listeners. The same trick, with smaller components, reaches the far higher frequencies of FM and television.

3. The choke in a fluorescent lamp

A fluorescent tube, once its gas has struck, behaves almost like a short circuit, and on the bare mains it would draw a runaway current and destroy itself in an instant. In series with it sits a **choke**, a coil of high inductance, whose job is to hold the current down. Because the choke opposes the current by its **reactance** $X_L = 2\pi fL$ rather than by resistance, it limits the current while dissipating almost no power of its own, which a plain resistor of the same limiting value never could.

At the 50Hz mains a choke of around 1H presents a reactance of some 300Ω, enough to settle the tube at its working current. This is the same ballast that lights the corridors of Miono Secondary School, and it is the inductor of the pure-circuit and series sections doing an honest day's work. Modern tubes replace the heavy iron choke with a small electronic ballast that does the same limiting at a far higher frequency.

4. Power-factor correction in industry

A workshop full of induction motors is a strongly inductive load: its current lags the voltage, its power factor is poor, and it therefore draws far more current than its real power alone would need. A bank of capacitors connected in parallel with the load supplies a leading current that cancels much of the motors' lagging reactive current, so the net reactive power falls, the power factor climbs back toward one, and the current the supply must carry drops, all without changing the useful work the motors do.

A Dar es Salaam metal workshop drawing 50kW at a power factor of 0.70 needs, as the design example of the power section found, a capacitor bank of about 640µF to reach a power factor of 0.95, trimming its line current from 172A to 127A. Because TANESCO bills industrial users on the apparent power in kVA, the bank pays for itself in lower demand charges, which is why capacitor banks hum quietly in the corner of almost every Tanzanian factory.

5. Induction heating and the induction cooker

In the transformer section the eddy currents induced in a solid core were a nuisance, fought off with laminations. Turn the same effect up deliberately and it becomes a tool. A coil carrying a high-frequency alternating current sets up a rapidly changing magnetic flux; a metal pan or workpiece placed in that flux has strong eddy currents driven round inside it, and its own resistance turns those currents straight into heat. The metal heats itself, from within, with no flame and no hot plate.

An induction cooker runs its coil at some 20kHz to 50kHz, heating only a ferromagnetic pot set on top while the glass surface beneath stays cool. On a larger scale, induction furnaces melt tonnes of steel by the same principle, and a smith's induction heater softens a bar in seconds. The loss the transformer designer works to suppress is, in the cook's hands, the whole point.

6. High-voltage supplies for hospital imaging

An X-ray tube cannot work on the mains as it stands: it needs tens of thousands of volts of steady direct current to fling electrons hard enough into its target to make X-rays. The supply that feeds it is the chapter in miniature. A step-up transformer raises the mains voltage enormously, exploiting the very turns ratio that the transformer section built; a rectifier and smoothing stage then convert that high alternating voltage into the steady high-voltage d.c. the tube demands.

The X-ray and CT equipment at Bugando Hospital runs its tubes at roughly 50kV to 120kV, voltages reached by stepping the mains up several hundredfold and then rectifying. The transformer that magnifies the voltage and the diodes that straighten it are exactly the devices of the previous two sections, here trusted with a diagnosis. The same arrangement, less dramatically, hides inside every microwave oven.

7. The metal detector

A metal detector is, at heart, a resonant LC circuit kept gently oscillating. Its search coil is the inductor; bring a piece of metal close and the metal alters the coil's inductance L, which shifts the circuit's resonant frequency $f_0 = \frac{1}{2\pi\sqrt{LC}}$ by a small but detectable amount. The electronics listen for that shift in frequency and sound the alarm. The whole instrument is the resonance section turned into a search tool.

Hand-held detectors run their coils at a few kHz up to hundreds of kHz, the lower frequencies reaching deeper and the higher ones catching smaller targets. They sweep for buried weapons at the doors of public buildings, and, in the goldfields around Geita, they help prospectors find the metal that pays. A shift in f_0 of a fraction of one per cent is all the machine needs.

We have followed electrical energy from its generation as alternating current, through its transmission and behaviour in AC circuits, to its final conversion into the direct current that powers everyday electronics. The Great Current Debate that opened the chapter is now, quietly, settled. Alternating current won the world not because it is gentler or simpler, it is neither, but because it alone can be transformed, and so alone can travel.

Yet understanding a subject is one thing; applying it confidently is another. Before we close the chapter, let us put these ideas to work in a final set of miscellaneous worked examples, drawing together the principles, techniques, and problem-solving strategies developed throughout our study of alternating current.

MISCELLANEOUS WORKED EXAMPLES ON ALTERNATING CURRENT THEORY

Example 43

(a) Explain why the EMF induced in a coil rotating at a steady speed in a uniform magnetic field varies sinusoidally with time, and state the position of the coil at which the induced EMF is (i) a maximum and (ii) zero.

(b) A rectangular coil of 50 turns, each of area $1.2 \times 10^{-2}\text{m}^2$, rotates at 3000 revolutions per minute in a uniform magnetic field of flux density 0.25T. Find (i) the peak EMF generated and (ii) the rms EMF.

Solution

(a) As the coil rotates, the flux linkage through it is $\Phi = \text{NBA}\cos\theta$, where θ is the angle between the field and the normal to the coil, and since the coil turns steadily $\theta = \omega t$. The induced EMF is the rate of change of this flux linkage, $E = -\frac{d\Phi}{dt} = \text{NBA}\omega\sin\omega t$, which is sinusoidal because the rate at which the flux is cut itself varies sinusoidally as the coil turns. The EMF is therefore a maximum when the plane of the coil is parallel to the field, since the sides then cut the lines of force at the greatest rate; and it is zero when the plane of the coil is perpendicular to the field, since the sides then move along the lines of force and cut none.

(b) The frequency is $f = \frac{3000}{60} = 50\text{Hz}$, and the angular speed is then:

$$\omega = 2\pi f = 2\pi \times 50\text{Hz} = 314\text{rad s}^{-1}$$

The peak EMF is:

$$E_0 = \text{NBA}\omega = 50 \times 0.25\text{T} \times 1.2 \times 10^{-2}\text{m}^2 \times 314\text{rad s}^{-1} = 47.1\text{V}$$

and the rms EMF is:

$$E_{\text{rms}} = \frac{E_0}{\sqrt{2}} = \frac{47.1\text{V}}{\sqrt{2}} = 33.3\text{V}$$

The coil generates a peak EMF of 47.1V and an rms EMF of 33.3V.

Example 44

(a) Explain what is meant by the root-mean-square (rms) value of an alternating current, and why it, rather than the peak or the mean value, is the value quoted for the mains supply and used in power calculations.

(b) An immersion heater connected to the 240V mains raises the temperature of 2.0kg of water from 20°C to 100°C in 7.0 minutes. Assuming no heat is lost to the surroundings, find (i) the power of the heater and (ii) its resistance. (Specific heat capacity of water $= 4200\text{J kg}^{-1}\text{K}^{-1}$.)

Solution

(a) The rms value of an alternating current is the value of the steady direct current that would dissipate heat in a given resistor at the same average rate as the alternating current does. It is defined in this way because the heating effect depends on the square of the current, so the average heating is governed by the mean of the square and not by the current itself; and since the alternating current spends as much time negative as positive, its ordinary average over a cycle is zero and would tell us nothing. The rms value is therefore the value quoted for the mains and used in power calculations, because it alone gives the true heating effect, so that $P = V_{\text{rms}}I_{\text{rms}}$ and the familiar d.c. formulae carry over unchanged.

(b) The heat needed to warm the water is:

$$Q = mc\Delta\theta = 2.0\text{kg} \times 4200\text{J kg}^{-1}\text{K}^{-1} \times (100 - 20)\text{K} = 6.72 \times 10^5\text{J}$$

Since all of it is supplied in $t = 7.0\text{min} = 420\text{s}$, the power is:

$$P = \frac{Q}{t} = \frac{6.72 \times 10^5 J}{420s} = 1600W$$

The resistance then follows from $P = \frac{V_{rms}^2}{R}$:

$$R = \frac{V_{rms}^2}{P} = \frac{(240V)^2}{1600W} = 36\Omega$$

The heater has a power of 1.6kW and a resistance of 36Ω.

Example 45

(a) The instantaneous alternating current $i = I_0 \sin\omega t$ and the displacement of a body in simple harmonic motion $x = A\sin\omega t$ are described by equations of exactly the same form. State the quantities in the a.c. case that play the parts of the amplitude A and the angular frequency ω of the S.H.M., and explain the sense in which a phasor is related to the circular motion that generates S.H.M.

(b) An alternating current is given by $i = 5.0\sin(100\pi t)A$, with t in seconds. Find (i) its frequency, (ii) its rms value, and (iii) the first instant after $t = 0$ at which the instantaneous current equals its rms value.

Solution

(a) In the a.c. equation the peak current I_0 plays the part of the amplitude A, since each is the greatest value its oscillating quantity reaches, while the angular frequency ω, related to the frequency by $\omega = 2\pi f$, plays the very same part in both. A phasor is related to the circular motion in just the way S.H.M. is: S.H.M. is the projection onto a diameter of a point moving steadily round a circle, and in the same way the instantaneous a.c. value is the projection of an arrow of length I_0 rotating steadily at angular frequency ω. So the phasor is that rotating arrow, and the alternating current is its shadow.

(b) Comparing with $i = I_0 \sin\omega t$ gives $I_0 = 5.0A$ and $\omega = 100\pi rad\ s^{-1}$.

(i) The frequency is:

$$f = \frac{\omega}{2\pi} = \frac{100\pi rad\ s^{-1}}{2\pi} = 50Hz$$

(ii) The rms value is:

$$I_{rms} = \frac{I_0}{\sqrt{2}} = \frac{5.0A}{\sqrt{2}} = 3.54A$$

(iii) The current equals its rms value when $\sin\omega t = \frac{1}{\sqrt{2}}$, that is when $\omega t = \frac{\pi}{4}$:

$$t = \frac{\pi/4}{\omega} = \frac{\pi/4}{100\pi rad\ s^{-1}} = \frac{1}{400}s = 2.5 \times 10^{-3}s$$

The current has frequency 50Hz, rms value 3.54A, and first reaches its rms value 2.5ms after $t = 0$.

Example 46

(a) A real coil, such as the field winding of a motor or the ballast of a fluorescent lamp, grows warm when it carries an alternating current, whereas an ideal (pure) inductor would stay cool. Explain why a real coil dissipates power while a pure inductor does not.

(b) A coil of resistance 30Ω and inductance 0.20H is connected in series with a 40Ω resistor across a 240V, 50Hz supply. Find (i) the impedance of the circuit, (ii) the current, (iii) the voltage across the coil, and (iv) the power factor.

Solution

(a) In a pure inductor the current lags the voltage by exactly a quarter-cycle, so that over each cycle the source pours energy into the magnetic field during one quarter and the field hands every joule back during the next; the average power is therefore zero and the component stays cool. A real coil, however, is wound from wire that has resistance as well as inductance, and this resistance carries the same current, so it dissipates energy at the rate I^2R as heat. The coil therefore grows warm because of its resistance and not its inductance, since only the resistive part of the coil draws real power from the supply.

(b) The inductive reactance is:

$$X_L = 2\pi fL = 2\pi \times 50Hz \times 0.20H = 62.8\Omega$$

The total resistance is $R = 30\Omega + 40\Omega = 70\Omega$, so the impedance is:

$$Z = \sqrt{R^2 + X_L^2} = \sqrt{(70\Omega)^2 + (62.8\Omega)^2} = 94.1\Omega$$

The current is:

$$I = \frac{V}{Z} = \frac{240V}{94.1\Omega} = 2.55A$$

The voltage across the coil is the current times the coil's own impedance, $Z_{coil} = \sqrt{(30\Omega)^2 + (62.8\Omega)^2} = 69.6\Omega$:

$$V_{coil} = IZ_{coil} = 2.55A \times 69.6\Omega = 178V$$

and the power factor of the whole circuit is:

$$\cos\phi = \frac{R}{Z} = \frac{70\Omega}{94.1\Omega} = 0.744 \text{ (lagging)}$$

The impedance is 94.1Ω, the current 2.55A, the voltage across the coil 178V, and the power factor 0.744 lagging.

Example 47

(a) In a loudspeaker system a capacitor is connected in series with the small high-frequency tweeter, and an inductor in series with the large low-frequency woofer. Explain, in terms of the way reactance depends on frequency, why this arrangement sends the high notes to the tweeter and the low notes to the woofer.

(b) A 0.10H inductor and a 50μF capacitor are available. (i) Find the reactance of each at 50Hz and state which offers the greater opposition. (ii) Find the frequency at which their reactances are equal.

Solution

(a) The reactance of a capacitor is $X_C = \frac{1}{2\pi fC}$, which falls as the frequency rises, so a capacitor offers little opposition to high notes but a large opposition to low ones; placed in series with the tweeter it therefore passes the high frequencies on to it while holding the low ones back. The reactance of an inductor is $X_L = 2\pi fL$, which rises with frequency, so an inductor offers little opposition to low notes but a large opposition to high ones; placed in series with the woofer it therefore passes the low frequencies to it while choking off the high ones. In this way each note is steered to the speaker that can reproduce it best.

(b) (i) At $f = 50\text{Hz}$ the reactances are:

$$X_L = 2\pi fL = 2\pi \times 50\text{Hz} \times 0.10\text{H} = 31.4\Omega$$

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50\text{Hz} \times 50 \times 10^{-6}\text{F}} = 63.7\Omega$$

So the capacitor offers the greater opposition at this frequency.

(ii) The reactances are equal when $X_L = X_C$, that is $2\pi fL = \frac{1}{2\pi fC}$, giving:

$$f = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{0.10\text{H} \times 50 \times 10^{-6}\text{F}}} = 71.2\text{Hz}$$

At 50Hz the capacitor dominates (63.7Ω against 31.4Ω), and the two reactances become equal at 71.2Hz.

Example 48

(a) In an audio amplifier a capacitor is connected in series between one stage and the next. Explain why such a capacitor allows the alternating audio signal to pass from one stage to the next, yet blocks the steady d.c. voltage used to bias each stage.

(b) A 10μF capacitor is connected across a 50V, 50Hz supply. Find (i) the rms current it draws, (ii) the maximum energy stored in it, and (iii) the maximum charge on its plates.

Solution

(a) No charge ever crosses the gap between a capacitor's plates; a current through a capacitor means only that the plates are being charged and discharged through the connecting wires. When the alternating signal is applied the voltage is forever changing, so the plates charge and discharge continually and a current flows in the wires at every moment, and the signal passes on. When the steady bias voltage is applied, however, the capacitor charges once to that voltage and then stops, because a constant voltage calls for no further movement of charge, so the current falls to zero and the d.c. is blocked. The capacitor therefore couples the signal while isolating the bias.

(b) (i) The capacitive reactance is:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50\text{Hz} \times 10 \times 10^{-6}\text{F}} = 318\Omega$$

Thus the rms current is:

$$I_{rms} = \frac{V_{rms}}{X_C} = \frac{50\text{V}}{318\Omega} = 0.157\text{A}$$

(ii) The maximum energy is stored at the peak voltage, $V_0 = V_{rms}\sqrt{2} = 50\text{V} \times \sqrt{2}$:

$$E_{max} = \frac{1}{2}CV_0^2 = \frac{1}{2} \times 10 \times 10^{-6}\text{F} \times \left(50\text{V}\sqrt{2}\right)^2 = 0.025\text{J}$$

(iii) The maximum charge is:

$$Q_{max} = CV_0 = 10 \times 10^{-6}\text{F} \times 50\text{V} \times \sqrt{2} = 7.07 \times 10^{-4}\text{C} = 707\mu\text{C}$$

The capacitor draws 0.157A, stores at most 0.025J, and holds at most 707μC.

Example 49

(a) Show that the average value of a sinusoidal alternating current over one complete cycle is zero.

(b) A half-wave rectified current of peak value 1.5A and frequency 50Hz is passed through a copper voltameter for 20 minutes. Find (i) the average (d.c.) current and (ii) the mass of copper deposited. (Electrochemical equivalent of copper $= 3.3 \times 10^{-7}\text{kg C}^{-1}$.)

Solution

(a) Let the current be $i = I_0 \sin\omega t$. Its average over one complete period T is the area under one cycle divided by the period:

$$I_{av} = \frac{1}{T}\int_0^T I_0 \sin\omega t \, dt = \frac{I_0}{\omega T}[-\cos\omega t]_0^T = \frac{I_0}{\omega T}(1 - \cos\omega T)$$

Since $\omega T = 2\pi$:

$$I_{av} = \frac{I_0}{\omega T}(1 - \cos 2\pi) = \frac{I_0}{\omega T}(1 - 1) = 0$$

The reason is plain in the graph: the current is positive for the first half-cycle and equally negative for the second, so the two areas cancel exactly and the mean over a whole cycle is zero.

(b) A half-wave rectifier removes the negative half-cycles and leaves only the positive humps, so the average over a full cycle is that of one hump spread over the whole period:

$$I_{av} = \frac{I_0}{\pi} = \frac{1.5A}{\pi} = 0.477A$$

The charge passed in $t = 20min = 1200s$ is:

$$Q = I_{av}t = 0.477A \times 1200s = 573C$$

By Faraday's first law the mass deposited is:

$$m = zQ = 3.3 \times 10^{-7} kg\, C^{-1} \times 573C = 1.89 \times 10^{-4} kg \approx 0.19g$$

The average current is 0.477A and about 0.19g of copper is deposited.

A pure alternating current, whose average is zero, would deposit nothing, which is why electroplating must use a rectified current.

Example 50

(a) Explain why a transformer will operate on alternating current but not on a steady direct current.
(b) An ideal step-up transformer has 200 turns on its primary and 5000 turns on its secondary. The primary is connected to the 240V, 50Hz mains and the secondary supplies a load of 1.0kW. Find (i) the secondary voltage, (ii) the secondary current, and (iii) the current drawn from the mains.

Solution

(a) A transformer works by mutual induction: the changing current in the primary sets up a changing magnetic flux in the core, and this changing flux, linking the secondary, induces an EMF in it. An alternating current is forever changing, so it produces a forever-changing magnetic flux and a steady alternating EMF appears across the secondary. A direct current, once established, is constant, so the flux it produces is constant too; and since an unchanging flux induces no EMF, the secondary gives nothing, apart from a brief pulse at the instant the current is switched on or off. A transformer therefore needs the perpetual change that only alternating current can provide.

(b) (i) For an ideal transformer $\frac{V_s}{V_p} = \frac{N_s}{N_p}$, so:

$$V_s = V_p \times \frac{N_s}{N_p} = 240V \times \frac{5000}{200} = 6000V$$

(ii) The secondary current follows from the load power, $P = V_s I_s$:

$$I_s = \frac{P}{V_s} = \frac{1000W}{6000V} = 0.167A$$

(iii) An ideal transformer wastes no power, so the input power equals the output power:

$$I_p = \frac{P}{V_p} = \frac{1000W}{240V} = 4.17A$$

The secondary delivers 6000V at 0.167A, while the primary draws 4.17A from the mains.

Example 51

(a) Name the main causes of energy loss in a practical transformer, and state one way in which each is reduced.
(b) A transformer supplies 4.0kW to its load while working at an efficiency of 96%. Its copper loss at this load is 90W. Find (i) the power drawn from the supply, (ii) the total power lost in the transformer, and (iii) the iron loss.

Solution

(a) There are three main losses.

- Copper loss is the I^2R heating in the windings, and it is reduced by winding them from thick, low-resistance copper.
- Eddy-current loss is the heating by currents induced in the core itself, and it is reduced by building the core from thin, insulated laminations, so that these currents are broken into small, weak loops.
- Hysteresis loss is the energy spent in repeatedly remagnetising the core each cycle, and it is reduced by making the core of soft iron, which has a narrow hysteresis loop and so wastes little energy per cycle.

(b) (i) Efficiency is $\eta = \frac{P_{out}}{P_{in}}$, so the input power is:

$$P_{in} = \frac{P_{out}}{\eta} = \frac{4000W}{0.96} = 4167W$$

(ii) The total loss is the difference between input and output:

$$P_{loss} = P_{in} - P_{out} = 4167W - 4000W = 167W$$

(iii) The iron loss is whatever remains after the copper loss is taken out:

$$P_{iron} = P_{loss} - P_{copper} = 167W - 90W = 77W$$

The transformer draws 4167W and loses 167W in all, of which 77W is iron loss.

Example 52

(a) Explain why connecting a capacitor in parallel with an inductive load, such as a motor, improves the power factor of the supply, even though the capacitor itself consumes no power.

(b) A single-phase motor takes 12A from the 415V, 50Hz mains at a power factor of 0.60 lagging. Find (i) the true power, (ii) the apparent power, (iii) the reactive power, and (iv) the capacitance that, connected in parallel, would raise the power factor to unity.

Solution

(a) An inductive motor draws a current that lags the voltage, and this lagging current carries a large reactive component which is drawn from and returned to the supply each cycle, doing no useful work yet still having to be carried by the line. A capacitor, by contrast, draws a current that leads the voltage, so its reactive current is in antiphase with the motor's. When the capacitor is connected in parallel, its leading reactive current cancels much of the motor's lagging reactive current, so the line need carry only the small difference. The total current therefore falls and comes more nearly into phase with the voltage, so the power factor rises, even though the capacitor itself takes no real power.

(b) (i) The true power is:

$$P = VI\cos\phi = 415V \times 12A \times 0.60 = 2988W \approx 2.99kW$$

(ii) The apparent power is:

$$S = VI = 415V \times 12A = 4980VA$$

(iii) With $\cos\phi = 0.60$, $\sin\phi = \sqrt{1 - 0.60^2} = 0.80$, so the reactive power is:

$$Q = VI\sin\phi = 415V \times 12A \times 0.80 = 3984VAR$$

(iv) For unity power factor the capacitor must supply all of this reactive power, $Q_C = V^2\omega C = Q$, so:

$$C = \frac{Q}{\omega V^2} = \frac{3984VAR}{2\pi \times 50Hz \times (415V)^2} = 7.4 \times 10^{-5}F = 74\mu F$$

The motor's true, apparent and reactive powers are 2.99kW, 4980VA and 3984VAR, and a 74μF capacitor restores unity power factor.

Example 53

(a) Explain, with reference to the action of the diodes, the difference between half-wave and full-wave rectification, and state which makes the better use of the alternating input.

(b) A sinusoidal supply of peak value 12V is applied, first through a single diode (half-wave) and then through a bridge of four diodes (full-wave), to a 100Ω load. Find the mean (d.c.) output voltage and the mean output current in each case, and hence the factor by which the full-wave circuit improves the d.c. output.

Solution

(a) In half-wave rectification a single diode conducts only during those half-cycles that drive it in the forward direction and blocks the reverse half-cycles, so only one half of each cycle reaches the load and the other half is wasted. In full-wave rectification a bridge of four diodes redirects the current so that both halves of every cycle pass through the load in the same direction, the reversed half-cycles being turned the right way up instead of being thrown away. Full-wave rectification therefore makes the better use of the input, since it gives twice the mean output and a waveform that is far easier to smooth.

(b) For half-wave rectification the mean voltage is $\frac{V_0}{\pi}$:

$$V_{dc} = \frac{V_0}{\pi} = \frac{12V}{\pi} = 3.82V$$

$$I_{dc} = \frac{V_{dc}}{R} = \frac{3.82V}{100\Omega} = 38.2mA$$

For full-wave rectification the mean voltage is twice as great, $\frac{2V_0}{\pi}$:

$$V_{dc} = \frac{2V_0}{\pi} = \frac{2 \times 12V}{\pi} = 7.64V$$

$$I_{dc} = \frac{7.64V}{100\Omega} = 76.4mA$$

The full-wave circuit gives 7.64V against 3.82V, an improvement of exactly a factor of two in the d.c. output.

Example 54

(a) Explain why a moving-coil meter reads zero when an alternating current is passed through it, whereas a hot-wire meter gives a reading, and state what value of the alternating current the hot-wire meter indicates.
(b) A current $i = (3.0 + 4.0\sin\omega t)A$, a steady current with an alternating part superimposed, flows through a circuit. Find (i) the reading of a moving-coil meter and (ii) the reading of a hot-wire meter placed in the circuit.

Solution

(a) A moving-coil meter gives a deflection proportional to the instantaneous current and in a direction that depends on the current's direction. When the current alternates, the deflecting torque reverses every half-cycle, and because the pointer cannot follow the rapid reversals it settles at the average value, which for an alternating current is zero. A hot-wire meter, by contrast, responds to the heating of a wire, and since the heating is proportional to i^2 it does not depend on the current's direction, so the wire is heated by both half-cycles alike. The hot-wire meter therefore reads the root-mean-square value of the current.

(b) (i) The moving-coil meter reads the average (d.c.) current. The average of the steady part is 3.0A and the average of $4.0\sin\omega t$ over a cycle is zero, so the meter reads 3.0A.

(ii) The hot-wire meter reads the rms value, the square root of the mean of i^2. Now:

$$i^2 = (9.0 + 24\sin\omega t + 16\sin^2\omega t)A^2,$$

Thus the mean square over a cycle is:

$$9.0A^2 + 0 + 16A^2 \times \frac{1}{2} = 17A^2$$

Hence:

$$I_{rms} = \sqrt{17A^2} = 4.12A$$

The moving-coil meter reads 3.0A, while the hot-wire meter reads 4.12A.

Example 55

(a) A tuned LC circuit lies at the heart of a radio's tuning dial and of a metal detector. Explain, in terms of energy, how an LC circuit that has been set oscillating behaves like a frictionless mass on a spring, and explain why in a real circuit the oscillations gradually die away.
(b) A capacitor of 8.0μF is charged to 2.0V and then connected across an inductor of 2.0mH of negligible resistance. Find (i) the maximum energy stored in the circuit, (ii) the maximum current in the inductor, (iii) the frequency of the oscillations, and (iv) the wavelength of an electromagnetic wave of this frequency.

Solution

(a) When the charged capacitor is connected to the inductor it drives a current that builds up a magnetic field, so the energy held in the capacitor's electric field is handed over to the inductor's magnetic field; and when that field collapses it recharges the capacitor the opposite way, so the energy shuttles endlessly between the two stores. This is exactly the exchange in a mass on a spring, where energy passes to and fro between the spring's potential energy (the capacitor) and the mass's kinetic energy (the inductor), the charge playing the part of displacement and the current that of velocity. In a real circuit the wire has resistance, so a little energy is turned to heat each cycle, just as friction drains a real pendulum; the oscillations are therefore damped and slowly die away.
(b) (i) At first all the energy is in the capacitor:

$$E_{max} = \frac{1}{2}CV^2 = \frac{1}{2} \times 8.0 \times 10^{-6}F \times (2.0V)^2 = 1.6 \times 10^{-5}J = 16\mu J$$

(ii) At maximum current all of this energy is in the inductor, $\frac{1}{2}LI_0^2 = E_{max}$, so:

$$I_0 = \sqrt{\frac{2E_{max}}{L}} = \sqrt{\frac{2 \times 1.6 \times 10^{-5}J}{2.0 \times 10^{-3}H}} = 0.13A$$

(iii) The frequency of free oscillation is:

$$f = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{2.0 \times 10^{-3}H \times 8.0 \times 10^{-6}F}} \approx 1.3 \times 10^3 Hz$$

(iv) The corresponding electromagnetic wavelength is:

$$\lambda = \frac{c}{f} = \frac{3.0 \times 10^8 m\ s^{-1}}{1.26 \times 10^3 Hz} = 2.4 \times 10^5 m$$

The circuit stores at most 16μJ and oscillates with a peak current of 0.13A at about 1.3kHz, corresponding to an electromagnetic wavelength of 2.4×10^5m, a reminder that radio tuning, at megahertz frequencies, needs far smaller L and C.

DIGGING DEEPER EXERCISE 13

EXERCISE 13A: BINDER QUESTIONS

Question 1

Explain why the mains current is said to reverse its direction one hundred times each second when its frequency is stated to be only 50Hz.

Question 2

The mains is described as a 240V supply, yet the insulation of the cables that carry it must be able to withstand a voltage of about 339V. Explain why.

Question 3

Explain, in terms of the way an inductor responds to a changing current, why the current through a pure inductor reaches its peak a quarter of a cycle later than the applied voltage.

Question 4

Explain why the opposition that a capacitor offers to an alternating current cannot be found by connecting an ordinary ohmmeter across it.

Question 5

A resistor of 30Ω is connected in series with an inductor whose reactance is 40Ω. Explain why the combined opposition to the current is not 70Ω.

Question 6

Explain why a series circuit containing a resistor, an inductor and a capacitor draws the greatest current from the supply at one particular frequency, and a smaller current at frequencies above and below it.

Question 7

Explain why a high quality factor is an advantage in the tuning circuit of a radio receiver, and state one disadvantage that would arise if the quality factor were made far too high.

Question 8

Explain why an inductor and a capacitor joined in parallel draw only a very small current from the supply at the resonant frequency, even though a large current flows back and forth within the loop itself.

Question 9

A voltmeter and an ammeter connected in an a.c. circuit read 240V and 5.0A. Explain why the power delivered to the circuit is not necessarily 1200W.

Question 10

Explain why the makers of a transformer or a generator quote its rating in kilovolt-amperes (kVA) rather than in kilowatts (kW).

Question 11

Explain why the two coils of a transformer are wound on a common iron core, and why the device would be very inefficient if the coils were merely placed side by side in air.

Question 12

Explain why the current drawn from the mains by the primary of a transformer increases when more appliances are switched on across its secondary.

Question 13

Explain why eddy currents are an unwanted source of energy loss in the core of a transformer, yet are the very effect on which the heating of a pan on an induction cooker depends.

EXERCISE 13B: REAL QUESTIONS

Question 14

A phone charger left plugged into a wall socket feels slightly warm even when no phone is connected to it. Explain why it continues to take power from the mains although it is charging nothing.

Question 15

Under the light of a fluorescent lamp running on the mains, the blades of a rotating fan sometimes appear to stand still, or to turn slowly backwards. Explain this observation.

Question 16

An electric motor, such as the one in a refrigerator or a water pump, draws a much heavier current at the instant it is switched on than when it has reached full running speed. Explain why.

Question 17

The large capacitor inside a television set or a power supply can give a serious electric shock when touched, even hours after the set has been switched off and unplugged. Explain why.

Question 18

When the switch controlling a large electromagnet or motor is opened, a bright spark jumps across the opening contacts, far brighter than the one seen when a plain lamp is switched off. Explain why.

Question 19

A simple electric wall clock driven from the mains keeps very accurate time over many months without ever being reset. Explain how the alternating supply makes this possible.

Question 20

A car battery can be recharged from a mains-operated charger, but not by connecting it to the secondary of a transformer alone. Explain why a rectifier must be included in the charger.

Question 21

A cheap multimeter measures the sinusoidal mains voltage correctly, but gives a noticeably wrong reading when used to measure the chopped output of a lamp dimmer. Explain why.

Question 22

A person who grips a wire carrying the 50Hz mains often cannot let go, whereas a brief shock from a direct-current source tends to throw the victim clear. Explain why.

Question 23

An electric arc welder is fitted with a transformer that has very many turns on its primary but only a few turns of very thick wire on its secondary. Explain why such a transformer is suited to welding.

Question 24

Lighting fitted in a bathroom or in a garden pond is often run at 12V from a transformer rather than directly at the 240V mains. Explain why this is done.

Question 25

During a power cut, mains appliances cannot be run straight from a 12V battery, but they work when the battery is connected through an inverter. Explain why the inverter is needed.

EXERCISE 13C: HOT QUESTIONS

Question 26

In a series circuit carrying an rms current of 0.50A, voltmeters read 60V across the resistor, 80V across the inductor and 50V across the capacitor. Find (a) the rms supply voltage, (b) the resistance, (c) the net reactance, (d) the phase angle between the supply voltage and the current, and (e) the average power dissipated.

Question 27

A coil draws an rms current of 2.0A from a 100V, 50Hz supply, and 1.4A from a 100V, 100Hz supply. Find the resistance and the inductance of the coil.

Question 28

An alternating current $i = 3.0\sin(100\pi t) + 4.0\sin(300\pi t)$, in amperes, flows through a 10Ω resistor. Find (a) the rms value of the current and (b) the average power dissipated in the resistor.

Question 29

A sinusoidal current of peak value 4.0A is passed through a half-wave rectifier and then through a 20Ω resistor. Find (a) the mean (d.c.) current, (b) the rms current, and (c) the average power dissipated in the resistor.

Question 30

A coil of inductance 0.20H and resistance 50Ω is connected in series with a 2.0μF capacitor across a 10V (rms), variable-frequency supply. Find (a) the resonant frequency, (b) the current at resonance, (c) the voltage across the capacitor at resonance, and (d) the quality factor of the circuit.

Question 31

A series tuned circuit is to resonate at 800kHz with a bandwidth of 8.0kHz, using a coil of inductance 0.20mH. Find (a) the quality factor, (b) the capacitance, and (c) the resistance required, and (d) state the bandwidth that would result if the resistance were doubled.

Question 32

A small power station generates 24kW at 480V. A transformer steps this up to 9.6kV for transmission along a cable of total resistance 6.0Ω. Find (a) the step-up turns ratio, (b) the current in the cable, (c) the power lost in the cable, and (d) the percentage of the generated power delivered to the far end.

Question 33

A 10kVA transformer has an iron loss of 80W and a full-load copper loss of 200W. For a load of power factor 0.80, find its efficiency (a) at full load and (b) at half load.

Question 34

A resistor of 50Ω, an inductor of reactance 25Ω and a capacitor of reactance 40Ω are connected in parallel across a 100V (rms) supply. Find (a) the current in each branch, (b) the line current and its phase angle, and (c) the reactance to which the capacitor must be changed to make the line current a minimum, and the value of that minimum current.

Question 35

A single-phase motor on the 240V, 50Hz mains takes a current of 8.0A while a wattmeter records 1.2kW. Find (a) the power factor, (b) the phase angle, (c) the reactive power, and (d) the extra current the supply must provide compared with a purely resistive load of the same true power.

Question 36

An inductor of inductance 0.50H carries an alternating current $i = 2.0\sin(100\pi t)$, in amperes. Find (a) the peak energy stored in the inductor, (b) the frequency at which the stored energy oscillates, (c) the rms current, and (d) the peak voltage across the inductor.

Question 37

A 100Ω resistor, a 0.30H inductor and a 20μF capacitor are connected in series across a 240V, 50Hz supply. Find (a) the impedance and current at 50Hz, (b) the power factor and average power, and (c) the frequency at which the power factor becomes unity, together with the current at that frequency.

Question 38

A transformer steps the 240V (rms) mains down to 9.0V (rms), whose output is full-wave rectified to charge a battery through a total resistance of 1.5Ω. Find (a) the turns ratio, (b) the peak secondary voltage, (c) the mean (d.c.) output voltage, and (d) the mean charging current.

Question 39

A coil of 500 turns, each of area 200cm², rotates at 3000 revolutions per minute in a uniform field of 0.040T and feeds a 100Ω resistor. Find (a) the peak EMF, (b) the rms EMF, (c) the rms current, and (d) the average power delivered to the resistor.

Question 40

A workshop draws 30kW at a power factor of 0.60 lagging from a 400V supply through a feeder of resistance 0.20Ω. Find (a) the line current and the power wasted in the feeder; then, after a capacitor bank raises the power factor to 0.90, find (b) the new line current and feeder loss, and (c) the percentage reduction in the feeder loss.

EXAMINATION QUESTIONS FOR PART THREE

Question 1

(a) Account for the fact that the electric field between two large, closely spaced, oppositely charged parallel plates is uniform, whereas the field around an isolated point charge grows weaker with distance from it.
(b) In a capacitor (condenser) microphone one plate is a light flexible diaphragm and the other is fixed, the pair being held at a constant charge. Account for how sound waves arriving at the diaphragm are turned into a varying electrical voltage.
(c) A parallel-plate capacitor has plates of area 150cm² separated by an air gap of 2.0mm and is charged to a potential difference of 200V. Calculate (i) the capacitance, (ii) the charge stored, (iii) the electric field between the plates, and (iv) the energy stored.

Question 2

(a) A charge Q is delivered to a capacitor that finally rests at a potential difference V. Account for the fact that the energy stored is one half of QV, and not QV.
(b) When a negatively charged rod is brought close to the cap of an uncharged gold-leaf electroscope, without touching it, the leaves diverge; they fall again when the rod is taken away. Explain this behaviour.
(c) A 4.0μF capacitor charged to 300V is connected in parallel with an uncharged 2.0μF capacitor. Calculate (i) the common potential difference, (ii) the final charge on each capacitor, and (iii) the energy lost in the process.

Question 3

(a) Account for the choice of infinity, rather than the charge itself, as the place of zero electric potential for a point charge.
(b) Explain why a small capacitor is often connected across a pair of switch contacts, in order to reduce the sparking seen when the switch is opened.
(c) A 50μF capacitor is charged to 100V and then discharged through a 10kΩ resistor. Calculate (i) the time constant, (ii) the initial discharge current, (iii) the charge remaining after one time constant, and (iv) the total heat dissipated in the resistor.

Question 4

(a) Account for the fact that, although the charge on a capacitor is proportional to the voltage across it, the ratio of the two, namely its capacitance, stays the same whatever the voltage is.
(b) Explain why a charged metal sphere loses its charge to the surrounding air much more rapidly on a humid day than on a dry one.
(c) An isolated metal sphere of radius 15cm stands in air, whose dielectric strength is 3.0×10^6V/m. Calculate (i) the capacitance of the sphere, (ii) the maximum potential to which it can be raised before the air breaks down, (iii) the corresponding charge, and (iv) the energy then stored.

Question 5

(a) Account for the fact that at the midpoint of the line joining two equal positive point charges the electric field is zero, yet the electric potential there is not zero.
(b) Explain why sliding a sheet of glass into the air gap of a charged parallel-plate capacitor increases its capacitance.
(c) A parallel-plate capacitor of plate area 200cm² and air gap 1.0mm is charged to 100V and then disconnected from the supply. A slab of mica of dielectric constant 5.0 is then slid in to fill the gap completely. Calculate (i) the capacitance before and after, (ii) the new potential difference, (iii) the field between the plates before and after, and (iv) the change in stored energy.

Question 6

(a) Account for the fact that electric field lines meet the surface of a charged conductor at right angles.
(b) A defibrillator stores energy in a capacitor and then discharges it through the patient. Explain why a capacitor is used for this, rather than connecting the patient straight to the charging supply.
(c) Three capacitors of 2.0μF, 3.0μF and 6.0μF are connected in series across a 90V supply. Calculate (i) the combined capacitance, (ii) the charge on each capacitor, and (iii) the voltage across each.

Question 7

(a) Account for the fact that the force between two point charges held a fixed distance apart is smaller when they are immersed in oil than when they are in air.
(b) Account for why an electrolytic capacitor can be made physically small yet have a very large capacitance compared with an air capacitor of the same size.
(c) Two small identical conducting spheres carry charges of +6.0nC and +2.0nC and are held 4.0cm apart in air. Calculate (i) the force between them; they are then momentarily joined by a fine wire and separated again; calculate (ii) the charge now on each and (iii) the new force between them.

Question 8

(a) Account for the fact that work must be done to pull the plates of a charged, isolated parallel-plate capacitor farther apart, and state what becomes of this work.
(b) Explain why the metal casing of a sensitive electronic instrument shields the circuit inside it from external electric fields.
(c) A parallel-plate capacitor has plates of area 100cm² at a separation of 1.0mm; it is charged to 50V and then disconnected. The plates are now pulled apart to a separation of 3.0mm. Calculate (i) the original capacitance and charge, (ii) the new capacitance, (iii) the new potential difference, and (iv) the work done in separating the plates.

Question 9

(a) Account for the fact that the electric field outside a uniformly charged sphere is the same as if all its charge were concentrated at its centre.
(b) Explain why the capacitor in a camera flash can light the bulb far more brightly, for a brief instant, than the small battery that charged it could on its own.
(c) A 2.0μF capacitor is charged to 200V. Calculate (i) the energy it stores; (ii) if it discharges fully in 2.0ms, the average power delivered; and (iii) state how this compares with the 5.0W that the charging battery can supply.

Question 10

(a) Account for the fact that a capacitor in a d.c. circuit passes a brief current when first connected but no steady current once it is fully charged.
(b) The dome of a Van de Graaff generator charged to several hundred thousand volts gives only a harmless tingle, yet the 240V mains can be lethal. Account for this.
(c) A 2.0μF capacitor is charged through a 1.5MΩ resistor from a 9.0V supply. Calculate (i) the time constant, (ii) the maximum charge, (iii) the charge after 3.0s, and (iv) the charging current at 3.0s.

Question 11

(a) Explain why a charged particle moving through a uniform electric field follows a parabolic path, just as a projectile does under gravity.
(b) Give reasons why every capacitor is marked with a maximum working voltage, and why exceeding it can permanently destroy the capacitor.
(c) An electron enters mid-way between two horizontal parallel plates 4.0cm long and 1.0cm apart, with 200V across them, moving horizontally at 4.0×10^{7}m/s. Calculate (i) the field between the plates, (ii) the vertical acceleration of the electron, (iii) its vertical deflection on leaving the plates, and (iv) the vertical velocity it has gained. (electron charge 1.6×10^{-19}C, mass 9.11×10^{-31}kg.)

Question 12

(a) Account for the fact that the electric field around a long straight charged wire falls off as 1/r with distance, whereas that of a point charge falls off as $1/r^2$.
(b) Explain why a coaxial cable, such as a television aerial lead, has capacitance between its inner and outer conductors, and why this capacitance matters when the cable carries high-frequency signals.
(c) A coaxial cable has an inner conductor of radius 1.0mm and an outer conductor of inner radius 4.0mm, with air between them, and is 50m long. Determine (i) its capacitance, and (ii) the charge stored when 200V is applied between the conductors.

Question 13

(a) Justify the view that the energy of a charged capacitor is stored in the electric field between its plates, rather than on the plates themselves.
(b) In electrostatic paint spraying the paint droplets are charged as they leave the nozzle. Explain how this makes the paint form an even coat on an earthed metal object, even reaching the surfaces that face away from the spray.
(c) A parallel-plate capacitor with plates of area 250cm² separated by 5.0mm of air is charged to 1.0kV. Calculate (i) the capacitance, (ii) the energy stored, and (iii) the energy density in the field, and verify that the energy density multiplied by the volume between the plates equals the stored energy.

Question 14

(a) Account for the fact that capacitors connected in series all carry the same charge, whereas capacitors connected in parallel all have the same voltage across them.
(b) Why is it that a plastic comb which has been run through dry hair can bend a thin stream of water falling from a tap, although the water carries no net charge?
(c) A 3.0μF capacitor is connected in series with a parallel combination of a 2.0μF and a 4.0μF capacitor, across a 60V supply. Find (i) the equivalent capacitance, (ii) the charge drawn from the supply, and (iii) the charge on, and the voltage across, the 2.0μF and 4.0μF capacitors.

Question 15

(a) Explain why a good capacitor dielectric should possess both a high dielectric constant and a high dielectric strength.
(b) Account for the fitting of sharp pointed conductors, called static wicks, on the trailing edges of an aircraft's wings and tail.
(c) Point charges of +2.0μC, +2.0μC and −3.0μC are placed at the corners of an equilateral triangle of side 10cm. Calculate the total electric potential energy of the configuration.

Question 16

(a) Give reasons why the whole of the charge given to a solid metal conductor resides on its outer surface.
(b) Explain why a supercapacitor can store far more charge than an ordinary capacitor of the same size, and why it is increasingly used to deliver short bursts of power.
(c) A 500μF capacitor is charged to 80V and then discharged through a 2.0kΩ resistor. Determine (i) the time constant, (ii) the initial discharge current, (iii) the energy dissipated in the resistor, and (iv) the time for the charge to fall to half its initial value.

Question 17

(a) What does the inverse-square law say about the force between two point charges? Give one experimental observation that supports it.
(b) Explain how the pressing of a key on a capacitive computer keyboard is detected electrically, without any mechanical contact being made.
(c) A small sphere of mass 0.20g carrying a charge of +5.0nC hangs by a light thread between two vertical parallel plates 5.0cm apart with 2.0kV across them. Calculate (i) the field between the plates, (ii) the horizontal force on the sphere, and (iii) the angle the thread makes with the vertical. ($g = 9.8 m/s^2$.)

Question 18

(a) Using Gauss's law, show that the electric field just outside a charged conducting surface is σ/ε_0.
(b) Account for the practice of immersing high-voltage capacitors in oil.
(c) The plates of a parallel-plate capacitor each carry a surface charge density of $2.0 \mu C/m^2$, have area $100 cm^2$ and are 2.0mm apart in air. Determine (i) the field between the plates, (ii) the potential difference, (iii) the capacitance, and (iv) the charge on each plate.

Question 19

(a) Explain why, when a small charged sphere is connected by a wire to a large uncharged one, most of the charge passes to the large sphere.
(b) Give reasons why the terminals of high-voltage equipment are made large and well rounded rather than small and sharp.
(c) A sphere of radius 5.0cm is charged to a potential of $9.0 \times 10^4 V$ and then connected by a long thin wire to a distant uncharged sphere of radius 10cm. Calculate (i) the initial charge on the first sphere, (ii) their common potential after connection, and (iii) the final charge on each sphere.

Question 20

(a) Account for the fact that the energy stored in a capacitor increases fourfold when the voltage across it is doubled.
(b) Explain why clothes made of synthetic fibres cling to the body and crackle when they are taken off in dry weather.
(c) A 10μF capacitor charged to 100V is discharged through a small electric motor that raises a 50g mass. If 60% of the energy stored in the capacitor becomes gravitational potential energy of the mass, calculate the height through which the mass is raised. ($g = 9.8 m/s^2$.)

Question 21

(a) Explain why the power dissipated in a resistor can be found from any of the expressions $P = VI$, $P = I^2R$ or $P = V^2/R$, and state which is the most convenient when (i) the current and resistance are known, and (ii) the voltage and resistance are known.
(b) Give reasons why the heating element of an electric kettle delivers the same heat whether it is run on a 12V steady (d.c.) supply or on a 12V (rms) alternating supply.
(c) A lamp marked "60W, 240V" is connected to the 240V, 50Hz mains. Calculate (i) its operating resistance, (ii) the rms current, (iii) the peak current, and (iv) the peak instantaneous power, and state the frequency at which the power delivered to the lamp fluctuates.

Question 22

(a) Account for the fact that the free electrons in a wire carrying alternating current merely oscillate back and forth about fixed positions and never travel from the source to the appliance.
(b) Explain why the connecting wires of a circuit carrying a large current stay cool, while the thin lamp filament carrying the same current becomes white-hot.

(c) A copper wire of cross-sectional area 2.0mm² carries an alternating current of peak value 4.0A at 50Hz. Taking the free-electron density as $8.5\times10^{28}/m^3$, determine (i) the peak drift velocity of the electrons and (ii) the amplitude of their back-and-forth oscillation. ($e = 1.6\times10^{-19}C$.)

Question 23

(a) Explain why the source used to charge a battery must have an electromotive force greater than that of the battery being charged.
(b) Account for the fact that a laptop charger contains not only a transformer but also a rectifier and a smoothing capacitor.
(c) The full-wave rectified output of a transformer provides a steady 14V, which charges a 12V battery of internal resistance 0.30Ω through a 1.2Ω protective resistor. Calculate (i) the charging current, (ii) the potential difference across the battery terminals, and (iii) the rate at which energy is stored chemically in the battery.

Question 24

(a) Account for the fact that a coil of wire offers a much greater opposition to alternating current than to direct current at the same voltage.
(b) Give reasons why an ammeter is made with a very low resistance and connected in series, whereas a voltmeter is made with a very high resistance and connected in parallel.
(c) An iron-cored coil of resistance 30Ω draws a current of 1.0A when connected to a 60V, 50Hz a.c. supply. Determine (i) its impedance, (ii) its inductive reactance, (iii) its inductance, and (iv) the power it dissipates.

Question 25

(a) Give reasons why an inductive load draws a larger current than a purely resistive load that delivers the same useful power at the same voltage.
(b) Explain why thick cables are used for the wiring that carries large currents, such as the supply to an electric cooker, while thin wire suffices for a lamp.
(c) A single-phase motor takes 6.0A at a power factor of 0.75 from the 240V, 50Hz mains, the supply cable having a resistance of 0.40Ω. Calculate (i) the true power, (ii) the apparent power, (iii) the current that a purely resistive load of the same true power would draw, and (iv) the extra power wasted in the cable owing to the motor's lower power factor.

Question 26

(a) Explain why a capacitor passes alternating current but blocks direct current, whereas an inductor passes direct current freely but opposes alternating current.
(b) Account for the use of a tightly coiled tungsten filament, sealed in an inert gas or a vacuum, in an electric lamp.
(c) A 10μF capacitor and a 2.0kΩ resistor are connected in series across a 12V, 50Hz a.c. supply. Find (i) the capacitive reactance, (ii) the impedance, (iii) the rms current, and (iv) the current that would flow if the same 12V were applied as direct current.

Question 27

(a) State, with a reason, why the value marked on a fuse and quoted for the mains supply is the rms value of the current or voltage, and not the peak value.
(b) Explain why running an electric heater for an hour costs far more than running a television for the same time on the same mains.
(c) A 240V (rms), 50Hz socket protected by a 13A fuse supplies an electric heater of resistance 24Ω. Calculate (i) the rms current drawn, (ii) the peak current, (iii) whether the fuse blows, and (iv) the energy used in 2.0 hours, in kilowatt-hours, and its cost at 290 TZS per kWh.

Question 28

(a) Account for the fact that, when resistors are joined in series, the same current passes through each, whereas when they are joined in parallel the same voltage acts across each.
(b) Explain why a transformer hums with a sound whose pitch is twice the frequency of the mains supplying it.
(c) A 100Ω resistor and a 200Ω resistor connected in parallel are placed in series with a coil of inductance 0.30H across a 240V, 50Hz supply. Determine (i) the resistance of the parallel combination, (ii) the inductive reactance of the coil, (iii) the impedance of the circuit, and (iv) the current drawn from the supply.

Question 29

(a) Explain why a wire carrying a current remains electrically neutral, even though it contains a stream of moving electrons.
(b) Account for the fact that a filament lamp run on the mains shows no visible flicker, although the current through it falls to zero a hundred times every second.
(c) An a.c. generator of rms EMF 250V and internal resistance 1.0Ω supplies a heater of resistance 49Ω. Find (i) the rms current, (ii) the terminal voltage, (iii) the power delivered to the heater, and (iv) the power wasted inside the generator.

Question 30

(a) Give two reasons why alternating current is preferred to direct current for the generation and distribution of electrical energy.
(b) Explain what is meant by the kilowatt-hour, and why electrical energy is sold in this unit rather than in joules.
(c) A 240V, 50Hz supply feeds the primary of an ideal step-down transformer of turns ratio 20:1, whose secondary drives a heating coil of resistance 0.60Ω. Calculate (i) the secondary voltage, (ii) the secondary current, (iii) the power delivered to the coil, (iv) the current drawn from the mains, and (v) the energy used in 5.0 hours, in kilowatt-hours.

Question 31

(a) Explain why the current is the same at all points in a simple series circuit, even when the circuit contains components of very different resistance.
(b) Account for the fact that the reading of a hot-wire a.c. ammeter does not change when the frequency of the supply is altered, provided the rms current stays the same.
(c) A 60Ω resistor and a 30Ω resistor are connected in parallel across a 240V, 50Hz supply. Calculate (i) the combined resistance, (ii) the rms current from the supply, (iii) the peak current, and (iv) the total power dissipated.

Question 32

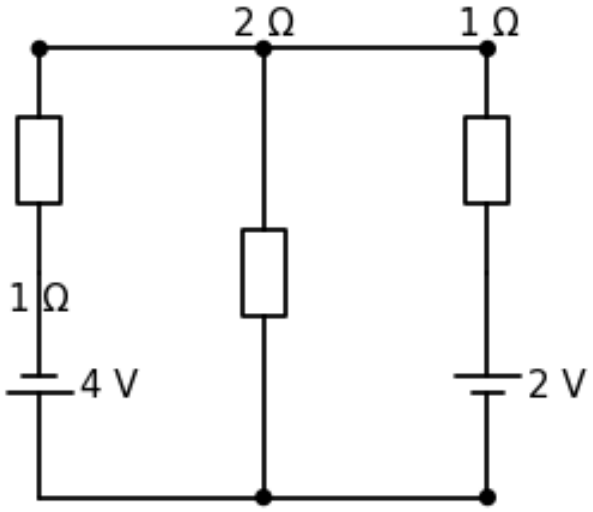
(a) Account for the fact that a pure inductor or a pure capacitor consumes no power on average, whereas a resistor does.
(b) Give reasons why a car battery can deliver a very large current to the starter motor, while a torch battery cannot.
(c) A pure inductor of reactance 60Ω is connected to a 240V, 50Hz supply, first by itself and then in series with an 80Ω resistor. Calculate, for each case, the current drawn and the average power dissipated.

Question 33

(a) Account for the fact that a potentiometer can measure the EMF of a cell accurately, whereas a voltmeter connected across the cell cannot.
(b) Explain why a choke (an iron-cored coil), rather than a resistor, is used to limit the current in a fluorescent-lamp fitting, and the advantage this gives in running cost.
(c) A 40Ω heating element is run from the 240V, 50Hz mains in series with a choke of inductance 0.15H and negligible resistance. Calculate (i) the reactance of the choke, (ii) the impedance, (iii) the current, and (iv) the power now dissipated in the heater, and compare it with the power dissipated without the choke.

Question 34

(a) Explain why, in a series a.c. circuit, the instantaneous voltages across the components always add up to the instantaneous supply voltage, while their rms (or peak) values do not.
(b) Account for the fact that, when two batteries of unequal EMF are connected in parallel, the stronger one may drive a current backwards through the weaker one.
(c) The figure shows two batteries and three resistors connected in two loops: a battery of EMF 4.0V in series with a 1.0Ω resistor forms the left branch, a battery of EMF 2.0V in series with a 1.0Ω resistor forms the right branch, and a 2.0Ω resistor forms the central branch common to both loops. Using Kirchhoff's laws, determine the current in each of the three branches.



Question 35

(a) Account for the fact that the primary of a transformer draws only a very small current when nothing is connected to its secondary.

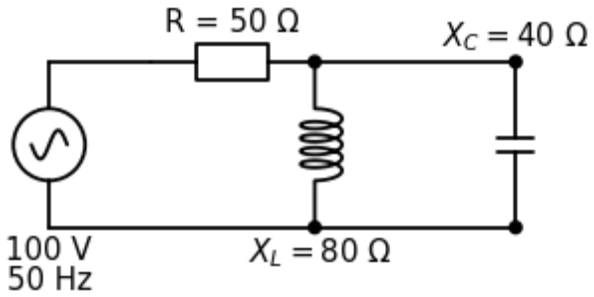
(b) Explain why the cables carrying current to a distant load must be made thick, or the load will receive a reduced voltage.
(c) A transformer supplies 5.0kW at 250V to a feeder cable of resistance 0.20Ω that carries the power to a load. Calculate (i) the current in the feeder, (ii) the voltage drop along it, (iii) the voltage at the load, and (iv) the power wasted in the feeder.

Question 36

(a) Give reasons why a fuse is rated by the current it can carry rather than by the voltage across it.
(b) Explain why an inductor and a capacitor of equal reactance, carrying the same current, have voltages across them that are exactly out of step.
(c) A circuit draws a current of 5.0A at a power factor of 0.80 lagging from the 240V, 50Hz mains. Determine (i) the true power, (ii) the reactive power, (iii) the impedance, and (iv) the resistance and the reactance of the circuit.

Question 37

(a) Explain why the impedance of a circuit that contains both inductance and capacitance depends on the frequency of the supply.
(b) Give reasons why the lights in a house may dim for a moment when a high-power appliance is switched on, if the supply wiring is thin.
(c) In the circuit shown, a 50Ω resistor is connected in series with a parallel combination of a pure inductor of reactance 80Ω and a pure capacitor of reactance 40Ω, across a 100V, 50Hz supply. Determine (i) the combined reactance of the parallel section, (ii) the impedance of the whole circuit, (iii) the current drawn from the supply, and (iv) the voltage across the parallel section.



Question 38

(a) Give reasons why connecting several cells in series increases the EMF available to a circuit.
(b) Explain why the reactance of an inductor is zero for direct current.
(c) A transformer rated 2.0kVA, 240V/24V has a full-load copper loss of 40W and an iron loss of 30W. At full load and a power factor of 0.90, calculate (i) the secondary current, (ii) the output power, (iii) the input power, and (iv) the efficiency.

Question 39

(a) Account for the fact that a step-down transformer can provide a larger current than is drawn from the supply.
(b) Account for the fact that the electrical energy used by a household is measured in kilowatt-hours rather than in joules.
(c) An 80Ω resistor and a coil of inductance 0.30H (of negligible resistance) are connected in series across a 240V, 50Hz supply. Find (i) the inductive reactance, (ii) the impedance, (iii) the current, (iv) the voltage across the resistor and across the coil, and show that these two voltages do not add arithmetically to 240V.

Question 40

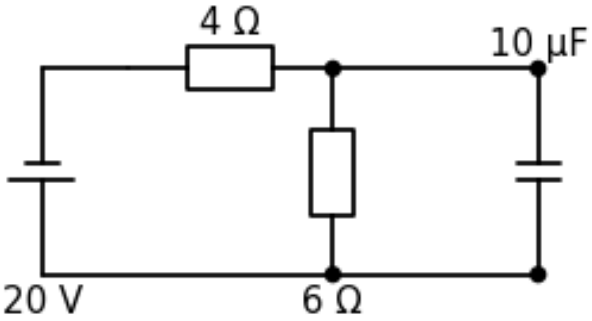
(a) Explain why the value of the unknown resistance found with a balanced Wheatstone bridge does not depend on the EMF of the supply used.
(b) Explain why an a.c. source delivers less average power to an inductive load than to a purely resistive load that draws the same current at the same voltage.
(c) The figure shows a Wheatstone bridge with arms P = 10Ω, Q = 20Ω, R = 15Ω and S = 25Ω, a galvanometer of resistance 30Ω connecting the junctions B and D, and a 6.0V battery of negligible internal resistance across A and C. Determine (i) whether the bridge is balanced, and (ii) the current through the galvanometer.

Question 41

(a) Explain why a capacitor passes a larger alternating current as the frequency of the supply is increased.
(b) Account for the fact that a capacitor intended for use on the 240V mains is marked with a working voltage of 400V or more.
(c) A 2.0μF capacitor is connected across a 240V, 50Hz supply. Calculate (i) its reactance, (ii) the rms current, (iii) the maximum charge on its plates, and (iv) the maximum energy stored in it.

Question 42

(a) Explain why, in the steady state, no current flows through a capacitor in a d.c. circuit, and why the voltage across it equals that across the component it is connected in parallel with.
(b) Account for the fact that a large capacitor connected across a supply helps to keep the voltage steady when the load current changes suddenly.
(c) In the circuit shown, a 4.0Ω resistor and a 6.0Ω resistor are connected in series across a 20V battery of negligible internal resistance, and a 10μF capacitor is connected across the 6.0Ω resistor. In the steady state, determine (i) the current in the resistors, (ii) the voltage across the 6.0Ω resistor, (iii) the charge on the capacitor, and (iv) the energy stored in it.



Question 43

(a) Explain why the electric field along a current-carrying wire is weak and almost uniform, unlike the strong field close to an isolated point charge.
(b) Give reasons why a thick wire carries a larger current than a thin wire of the same length and material when the same voltage is applied across it.
(c) A potential difference of 1.5V is applied across a copper wire 2.0m long and of cross-sectional area 1.0mm^2. (Resistivity of copper = $1.7\times10^{-8}\Omega$ m; free-electron density = 8.5×10^{28}m^{-3}.) Calculate (i) the resistance of the wire, (ii) the current, (iii) the electric field along the wire, and (iv) the drift velocity of the electrons.

Question 44

(a) Explain why a charged particle gains kinetic energy as it moves through a potential difference, and state where this energy comes from.
(b) Account for the fact that increasing the accelerating voltage in a cathode-ray tube produces a faster electron beam.
(c) An electron, starting from rest, is accelerated through a potential difference of 4.0kV across a gap 5.0cm wide. Calculate (i) the kinetic energy it gains (in joules and in electron-volts), (ii) its final speed, (iii) the uniform electric field in the gap, and (iv) the force on the electron. ($e = 1.6\times10^{-19}$C, $m_e = 9.11\times10^{-31}$kg.)

Question 45

(a) Explain why the output of a full-wave rectifier, although always in one direction, is not steady and must be smoothed.
(b) Account for the fact that a capacitor of larger capacitance smooths the rectified output more effectively.
(c) A transformer steps the mains down to a peak of 12V, which is full-wave rectified and then smoothed by a 2200μF reservoir capacitor across a 60Ω load. Determine (i) the mean (d.c.) output voltage before smoothing, (ii) the charge on the reservoir capacitor when charged to the peak, (iii) the energy then stored in it, and (iv) the mean load current before smoothing.

Question 46

(a) Account for the fact that a potential difference must be maintained across a conductor for a steady current to flow through it.
(b) Give reasons why the connecting wires of a circuit are usually taken to have no resistance.
(c) A 12V battery of internal resistance 0.50Ω drives a current through an external resistor of 5.5Ω, which is a uniform wire 10cm long. Find (i) the current, (ii) the terminal voltage of the battery, (iii) the electric field along the resistor wire, and (iv) the power dissipated in the resistor.

Question 47

(a) Explain why a charged capacitor can deliver a large current for a brief instant when it is discharged through a low resistance.
(b) Account for the fact that discharging a given amount of stored energy in a shorter time produces a larger average power.
(c) A 50μF capacitor charged to 100V is discharged completely through a small motor in 0.50s. Calculate (i) the charge that flows, (ii) the energy released, (iii) the average current, and (iv) the average power delivered.

Question 48

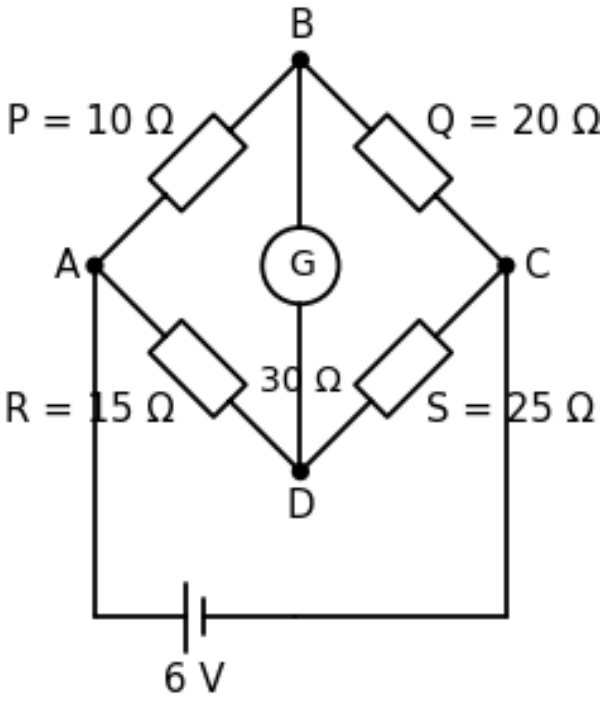
(a) Explain why a series combination of an inductor and a capacitor responds most strongly at one particular frequency.
(b) Account for the fact that, in an oscillating LC circuit, the energy is stored entirely in the capacitor at one instant and entirely in the inductor a quarter of a cycle later.
(c) A coil of inductance 0.40H is connected in series with a 10μF capacitor across a supply of variable frequency. Determine (i) the resonant frequency, and (ii) the reactance of the inductor and of the capacitor at that frequency.

Question 49

(a) Give reasons why connecting two capacitors in series gives a combined capacitance smaller than either of them.
(b) Explain why, once the capacitors in a d.c. circuit are fully charged, the whole of the battery's EMF appears across them.
(c) A 4.0μF capacitor and a 2.0μF capacitor are connected in series across a 30V battery. Find (i) the combined capacitance, (ii) the charge on each capacitor, and (iii) the voltage across each.

Question 50

(a) Account for the fact that a capacitor blocks a steady direct current but allows an alternating current to pass.
(b) Explain why a capacitor draws a current from an a.c. supply yet consumes no power on average.
(c) A 20μF capacitor is connected first to a 100V battery and then to a 100V, 50Hz a.c. supply. Calculate (i) the charge it holds and the steady current it passes on the battery, and (ii) its reactance, the rms current, and the average power on the a.c. supply.



ANSWERS TO DIGGING DEEPER EXERCISES

EXERCISE 1A

1. Hooke's law states that within the elastic regime, the extension of a wire is directly proportional to the applied force, so doubling the force doubles the extension. This proportionality holds because the interatomic restoring forces in the wire respond linearly to small displacements of the atoms from their equilibrium positions. Once the stress exceeds the ***proportional limit,*** the linearity breaks down because the interatomic forces are no longer well approximated as linear, and the extension begins to grow faster than the force.

2. The long thin wire extends more. The extension is given by $\Delta L = FL/(AE)$. For the same material (same E) and the same force (F), the extension is directly proportional to the original length L and inversely proportional to the cross-sectional area A. The long thin wire has a larger L and a smaller A, both of which act to increase the extension. The two geometric factors that determine the answer are therefore the original length and the cross-sectional area.

3. Young's modulus E describes resistance to longitudinal (tensile or compressive) deformation, and is the relevant property for a wire, rod, or column being stretched or compressed along its length. The shear modulus G describes resistance to angular (shape-changing) deformation at constant volume, and is the relevant property for a block being sheared sideways or a shaft being twisted. The bulk modulus B describes resistance to uniform volumetric (size-changing) deformation at constant shape, and is the relevant property for a fluid or solid subjected to a uniform hydrostatic pressure from all directions.

4. A fluid has no permanent intermolecular forces resisting tangential displacement of one layer relative to another. A static shear stress applied to a fluid causes its layers to flow continuously past each other, so the fluid never reaches a static equilibrium under shear and its shear modulus is zero. A fluid does, however, resist a uniform compression because squeezing the molecules closer together still meets a strong repulsive force, so its bulk modulus is finite. The essential physical difference between a solid and a fluid is therefore the presence in a solid of intermolecular forces that resist shape change, and their absence in a fluid.

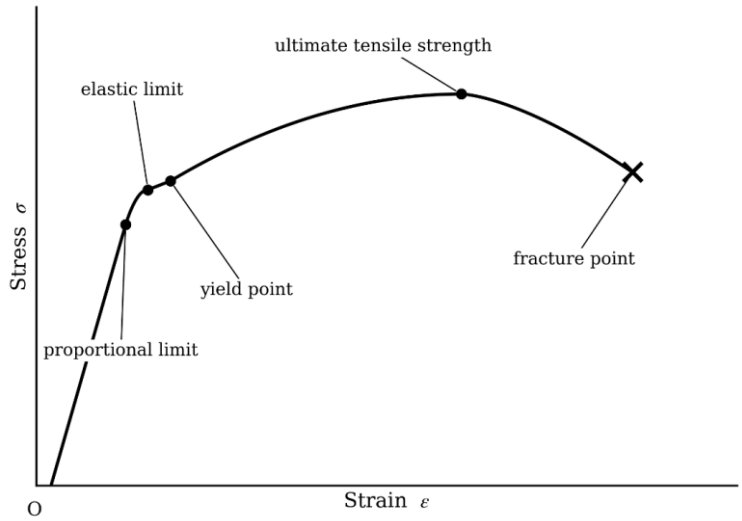
5. By Hooke's law the extension is $e = \frac{FL}{AE}$. Since F, L and A are the same for both, $e \propto \frac{1}{E}$; and because rubber has a far smaller Young's modulus E than steel, the rubber band extends more.

The elastic potential energy stored is $U = \frac{1}{2}Fe$, the area under the force-extension graph. The force F is the same, but the rubber has the greater extension e, so the rubber band also stores more elastic potential energy.

6. The total elastic potential energy stored in a stretched wire is the energy contained in the entire piece, measured in joules. The energy per unit volume is the energy stored in each cubic metre of the material, measured in joules per cubic metre. The total energy depends on the size of the wire, while the energy per unit volume depends only on the material and the stress applied. When comparing two different materials, the total energy is uninformative because it conflates the size of the sample with the property of the material. The energy per unit volume is the meaningful comparison because it isolates the material property from the geometry, and shows directly which material is intrinsically better at storing elastic energy.

7. Elastic fatigue is the gradual weakening of a material under many cycles of loading at stresses well below its breaking stress, caused by the slow growth of microscopic cracks at surface flaws with each loading cycle. It differs from ordinary overloading, which produces an immediate plastic deformation in a single event when the elastic limit is exceeded. Elastic fatigue accumulates damage cycle by cycle without any visible warning, while overloading damage is immediate and obvious. An everyday example is a spring balance in daily use at a market, whose zero reading drifts slowly over months and years because of elastic fatigue in its spring, even though no single load has ever come close to overloading it.

8.



- Between the origin and the proportional limit, the metal obeys Hooke's law and stress is directly proportional to strain.
- Between the proportional limit and the elastic limit, the metal still returns to its original length on unloading but the stress-strain relation is no longer linear.
- Between the elastic limit and the yield point, plastic flow begins and any unloading would leave a permanent set.
- Between the yield point and the ultimate tensile strength, the metal undergoes large plastic strain at gradually increasing stress as it work-hardens.
- Between the ultimate tensile strength and the fracture point, the metal forms a localised neck that thins rapidly until it breaks.

9. When a bimetallic strip is heated, both metals try to expand, but the metal with the larger coefficient of linear expansion tries to grow more than the metal with the smaller coefficient. The bond between them prevents either layer from sliding relative to the other, so the strip bends into an arc. The metal with the larger coefficient of expansion (brass, $\alpha \approx 19 \times 10^{-6}K^{-1}$) lies on the outer side of the curve, because the outside of a curve is longer than the inside, and the brass needs the extra length to accommodate its greater natural expansion. A practical device that depends

on this bending is the mechanical thermostat, in which the bimetallic strip opens or closes an electrical contact as the temperature crosses a set point.

10. Searle's two-wire apparatus uses two identical wires hanging side by side from the same rigid support. One wire is the test wire, which carries the variable load whose effect on its extension is to be measured. The other is the reference wire, which is kept under a constant small load just heavy enough to keep it straight. A light frame attached to each wire connects them through a sensitive spirit level fitted with a micrometer screw gauge. As the test wire stretches under the applied load, the spirit level tips slightly and the micrometer is adjusted to recentre the bubble. The reading gives the extension of the test wire, and Young's modulus is then extracted from the slope of the load-extension graph.

The two advantages of the two-wire arrangement over a single-wire experiment are: **(1) thermal compensation,** because both wires hang in the same air at the same temperature, so any expansion or contraction caused by room temperature changes affects both equally and cancels out of the difference; **(2) compensation for any small movements of the support beam,** because both wires are suspended from the same beam and any settling, vibration, or unintended motion of the beam affects the tops of both wires equally and does not appear in the difference.

EXERCISE 1B

11. The spring is suffering from elastic fatigue. Each time Mama Amina hangs a load on the balance and then removes it, the spring is stretched and released within its elastic limit, but microscopic flaws on the surface of the wire grow a tiny amount with each cycle. After two years and many thousands of weighings, this damage has accumulated into a small permanent plastic deformation of the spring, which is why the unloaded length is now longer than when it was new and the pointer reads 50g instead of zero. Applying a correction of -50g to every reading is not a complete solution because the elastic fatigue continues to accumulate with every new use, so the zero error will keep growing. Worse, **fatigue can also reduce the spring constant** of the spring, so the same load no longer produces the same extension, meaning that the entire calibration has shifted and a single zero correction will not restore accuracy across the full range of the balance.

12. A long bridge in a coastal city experiences large temperature swings between cold nights and hot afternoons. If the steel cables and deck were rigidly fixed at both ends, the prevented thermal expansion would generate a thermal stress $\sigma = E\alpha\Delta T$ large enough to crack the anchor blocks and fatigue the steel. Expansion joints allow the bridge to lengthen and shorten freely with temperature, so that the temperature change produces motion rather than stress.

13. The tyre acts as an elastic shock absorber between the road and the rigid metal frame of the car. Every small bump, pothole, or stone in the road would otherwise transmit a sharp impulsive force directly to the wheel rim and from there to the suspension and the passengers. The compliant rubber and air-filled structure of the tyre deforms elastically over a small distance when it meets a bump, spreading the impact force over a longer time and converting much of the impulsive shock into a gentle compression. A car with completely rigid wheels would transmit every roughness in the road surface, however small, as a sharp jolt to the chassis. The ride would be intolerably rough, the suspension would wear out quickly from the high-frequency loading, and the metal frame and bodywork would be subjected to repeated impulsive stresses that would fatigue them rapidly.

14. When the truck crosses the riveted bridge, the steel members bend slightly under the changing load, and at each riveted joint the plates slip a microscopic distance against one another before friction holds them again. Each microscopic slip releases a small burst of stored elastic energy as sound, and the accumulated bursts from many joints across the bridge produce the creaking. A welded bridge is silent because welded joints are continuous metal with no slipping surfaces, so there is no friction-driven release of elastic energy at the junctions.

15. When a horizontal beam bends under a downward load, the top of the beam is in compression and the bottom is in tension. Concrete is strong in compression but weak in tension, so a pure concrete beam would crack along the bottom. The steel bars are placed near the bottom because that is where the tensile stress is largest, and steel has a tensile strength which by far is higher than that of concrete; the steel therefore carries the tension while the concrete carries the compression.

16. The hard outer shell stops almost instantly on impact and would transmit the full deceleration directly to the rider's head if there were nothing soft inside it. The crushable foam liner deforms plastically over several centimetres during the impact, extending the stopping distance of the head and converting the kinetic energy of the head into the work of crushing the foam. Because the stopping distance is much longer, the peak deceleration on the brain is much smaller, and the impact is survivable.

17. The crack started at the corner because corners are points of high stress concentration: any uniform stress applied across a glass panel is amplified at sharp corners by a large factor (just as the corners of the de Havilland Comet windows became fatigue crack starters). The missing rubber spacers were responsible for producing the stress in the first place. Without rubber spacers to absorb the thermal expansion of the glass relative to the wooden frame, the glass developed a thermal stress $\sigma = E\alpha\Delta T$ each time the room warmed in the afternoon, and this stress was amplified at the corners until a crack opened.

18. A sudden jerk extends the stiff nylon line through a very small distance before it reaches its breaking point, which means the line absorbs the kinetic energy of the rod and the fish over a tiny stopping distance. From $U = \frac{1}{2}F\Delta L$, a small ΔL requires a very large peak force F to absorb the same energy, and that peak force exceeds the breaking strength even though the average force during a steady pull would not. A stretchy line would have spread the same energy over a longer extension and reduced the peak force enough to keep the line intact.

19. Strength and stiffness are independent properties of a material. Old bone has roughly the same Young's modulus as young bone (so it feels just as stiff under small loads) but a much smaller breaking stress, because the porous microstructure that develops with age provides easier paths for cracks to propagate. The grandmother's bone was therefore stiff enough to feel healthy right up to the moment of the fall, then failed at a stress that a young bone would have absorbed without difficulty.

EXERCISE 1C

20. Deceleration: $a = \frac{\Delta v}{\Delta t} = \frac{3.0}{0.40} = 7.5\text{m/s}^2$ (upward).

Net upward force on cage: $F_{net} = ma = (1500)(7.5) = 11250\text{N}$.

Tension in cable: $T = mg + ma = 1500(9.8 + 7.5) = 25950\text{N}$.

$\sigma = \frac{T}{A} = 25950/(4.0 \times 10^{-4}) = 6.49 \times 10^{7}\text{Pa}$

21. $125°\text{C}$

22. $8.03 \times 10^{8}\text{Pa}$

23. 17.7m/s

24. $U_A : U_B = 3:5$

25. $2.58 \times 10^{6}\text{N}$

26. 25.09mm.

27. $228°\text{C}$

28. Maximum working stress: $\sigma_{working} = \frac{\sigma_{yield}}{n} = \frac{2.5\times10^{8}}{5} = 5.0\times10^{7}\text{Pa}$.

Area of one cable: $A = \frac{\pi(5.0\times10^{-3})^2}{4} = 1.963\times10^{-5}\text{m}^2$.

Maximum force per cable: $F_{max} = \sigma_{working}A = (5.0\times10^{7})(1.963\times10^{-5}) = 981.7\text{N}$.

Total weight allowed: $W_{max} = 6F_{max} = 5890\text{N}$.

$$m_{max} = \frac{W_{max}}{g} = \frac{5890}{9.8} = 601\text{kg}$$

29. 3.83kg

30. For a conical pendulum, the vertical component of the wire tension supports the weight: $T\cos30° = mg$

$$T = \frac{mg}{\cos30°} = \frac{(0.20)(9.8)}{0.866} = 2.263\text{N}$$

Cross-sectional area: $A = \frac{\pi(0.50\times10^{-3})^2}{4} = 1.963\times10^{-7}\text{m}^2$

$$\Delta L = \frac{TL}{AE_c} = \frac{(2.263)(0.80)}{(1.963\times10^{-7})(1.2\times10^{11})} = 7.69\times10^{-5}\text{m} \approx 0.077\text{mm}$$

31. Fractional density change is approximately equal to the negative of the fractional volume change:

$$\frac{\Delta\rho}{\rho_0} = \frac{(1003.0-1000.0)}{1000} = 3.0\times10^{-3}$$

From the bulk modulus, $\Delta P = -B\left(\frac{\Delta V}{V_0}\right) = B\left(\frac{\Delta\rho}{\rho_0}\right)$:

$$\Delta P = (2.2\times10^{9})(3.0\times10^{-3}) = 6.6\times10^{6}\text{Pa}$$

The pressure at depth h is $\Delta P = \rho_0 gh$:

$$h = \frac{\Delta P}{\rho_0 g} = \frac{(6.6\times10^{6})}{(1000)(9.8)} = 673\text{m}$$

32. Conservation of energy. As the ball falls a distance ΔL (the maximum extension of the wire), it loses gravitational PE equal to $mg\Delta L$ and gains elastic PE equal to $\frac{1}{2}k(\Delta L)^2$, where $k = AE/L$ is the spring constant of the wire:

$$mg\Delta L = \frac{1}{2}\frac{AE}{L}(\Delta L)^2$$

Solving for ΔL:

$$\Delta L = \frac{2mgL}{AE} = \frac{2(2.0)(9.8)(3.0)}{(1.0\times10^{-6})(2.0\times10^{11})} = 5.88\times10^{-4}\text{m} \approx 0.59\text{mm}$$

This is twice the static extension of the wire under the same weight, as expected for a sudden load applied to a Hookean wire.

33. $7.5\times10^{-5}\text{m}$ (or 0.075mm)

EXERCISE 2A

1. Kipanga's common sense is not common physics. The Zeroth Law guarantees that "being in thermal equilibrium" is a **transitive** relation: if A and B are in thermal equilibrium, and B and C are in thermal equilibrium, then A and C must also be. This is not a mathematical identity; it is an empirical law that has to be tested experimentally, because there is no a priori reason why thermal equilibrium should behave transitively. Without it, a thermometer placed in body B could not be used to compare body A with body C. The Zeroth Law is therefore the foundation on which every thermometer's logic rests.

2. The mistake is the claim that a body "contains" heat. Heat is not a quantity stored inside a body; it is energy transferred between two bodies because of a temperature difference between them. A body can have a temperature, and it can exchange heat with another body.

3. (a) A thermometric property is a physical property that changes **reproducibly and monotonically** with temperature and can therefore be used to measure temperature. On another hand, an ordinary physical property is any measurable feature of a substance, not all of which satisfy this criterion.

(b) Two examples of non-thermometric properties:

- The colour of a metal at room temperature, because it does not change with small temperature changes.
- The density of water between 0°C and 4°C, because it does not change monotonically with temperature (it increases then decreases), so a single density value corresponds to two temperatures.

4. Two disadvantages of the ice point:

- The ice point depends on the atmospheric pressure above it. A laboratory at high altitude with lower atmospheric pressure measures an ice point at a slightly different temperature from a laboratory at sea level.
- The ice point is affected by dissolved air in the water. Pure water with no dissolved air freezes at a slightly higher temperature than water saturated with dissolved air, and the difference is large enough to be significant at the precision required for a primary reference.

5. (a) Thermal equilibrium between the thermometer and the bath.

(b) Before thermal equilibrium is reached, the thermometer is at a different temperature from the bath, and its reading reflects only its own temperature, not the bath's. Only at thermal equilibrium does the thermometer's temperature equal the bath's, and only then does the reading correspond to the true temperature of the bath.

6. (a) Water has the largest value of specific heat capacity among common liquids. The numerical value is $4200\text{J kg}^{-1}\text{K}^{-1}$.

(b) A given mass of oil absorbs less than half the heat per degree of temperature rise compared with the same mass of water. For the same cooling duty, a system using oil would need either more than twice the mass flow rate of oil, or allow the oil to heat up by more than twice as much, either of which would require larger pumps, larger pipes, and larger radiators. Water's larger specific heat capacity makes it the more efficient coolant, whenever it is available.

7. During melting, intermolecular forces holding the solid rigid structure must be partially loosened to allow the molecules to slide past one another as a liquid. Loosening the intermolecular forces requires energy, which is supplied as heat absorbed from the surroundings. During freezing, the reverse process occurs: the molecules lock back into rigid positions, and the energy previously used to loosen the intermolecular forces is released

as heat to the surroundings. The temperature stays at 0°C throughout because the energy change concerns only the molecular potential energy (bond configuration), not the molecular kinetic energy (which sets the temperature).

8. (a) Three differences:

- Evaporation occurs only at the surface of a liquid; boiling occurs throughout the volume.
- Evaporation occurs at any temperature above the freezing point; boiling occurs only at a specific temperature (the boiling point) that depends on the applied pressure.
- Evaporation is a slow process; boiling is rapid and produces visible bubbles in the liquid.

(b) Both processes absorb the same latent heat of vapourisation per kilogram of liquid converted to vapour, because the molecular cost of separating a liquid molecule into the vapour phase is the same regardless of where the separation happens.

9. A thermometer placed in contact with a body extracts heat from the body to warm itself, and the amount extracted is determined by the thermometer's thermal capacity which is small for small bulb. If the thermometer's thermal capacity is very small compared with the body's, the heat extracted disturbs the body's temperature only **negligibly**, and the reading is close to the body's original temperature. If the thermometer's thermal capacity were comparable to the body's, the act of measurement would significantly cool the body, and the reading would correspond to the disturbed temperature rather than the original one. The measurement would then be meaningless. A small bulb minimises the mass of the thermometric substance and therefore the disturbance.

EXERCISE 2B

10. (a) Kipute is using temperature correctly. Temperature is a property of the air (or of a body), not a property of a person's response to it. Kipanga is inferring the air's temperature from his own sweating, but sweating also depends on what he ate, how much he is exerting himself, his hydration level, and his clothing. These personal factors vary from day to day, so sweating is not a reliable indicator of the air's temperature.

(b) A thermometer measures a physical property of a substance inside the thermometer itself such as the length of a mercury column, the resistance of a platinum coil, the EMF of a thermocouple junction, that changes reproducibly and only with temperature. The reading depends on the thermometer's physics and not on the person using it.

11. (a) Two reasons:

- The thermometer begins each measurement well below body temperature (the antiseptic bath is cool), so during the brief time in the mouth it absorbs heat rapidly but has not yet reached thermal equilibrium when she reads it.
- Shaking the thermometer immediately before insertion leaves no time for the mercury column to settle from its lowest position, and the reading reflects a transient state rather than the true equilibrium temperature.

(b) Two changes:

- Allow each thermometer to remain in the child's mouth long enough (typically two to three minutes) to reach thermal equilibrium before reading.
- Rinse and dry the thermometer between patients instead of cooling it in antiseptic, so it starts each measurement near room temperature rather than well below it.

12. The father's reasoning is wrong because the gas flame is not twice as hot as the charcoal flame; the two flame temperatures are broadly similar. What differs is the **rate at which each cooker delivers heat** to the pot. The gas flame is more concentrated and has a much higher heat transfer rate to the base of the pot, so more joules of heat per second reach the beans. The charcoal flame, being diffuse and partially shielded by ash, delivers heat more slowly. Since the heat required to cook the beans is the same in both cases, the faster delivery rate of the gas cooker means the same heat is delivered in a shorter time.

13. The water-filled bottle will keep Mama Juma warmer for longer. If both bottles lose heat to the surroundings at the same rate, the bottle that can release more heat per degree of its own temperature drop is the one that stays warm longer. The water's specific heat capacity is $4200\,\mathrm{J\,kg^{-1}K^{-1}}$, more than double the oil's $2000\,\mathrm{J\,kg^{-1}K^{-1}}$, so a kilogram of water at 60°C stores more than twice the heat (above the room temperature) that a kilogram of oil at 60°C stores. With the same rate of heat loss, the water bottle's temperature falls more slowly than the oil bottle's, and Mama Juma stays warm for roughly twice as long.

14. (a) Water has a specific heat capacity of $4200\,\mathrm{J\,kg^{-1}K^{-1}}$, higher than any other common liquid. A small mass of water flowing through the engine can absorb large quantities of heat for only a modest temperature rise. Used engine oil has a much smaller specific heat capacity, roughly $2000\,\mathrm{J\,kg^{-1}K^{-1}}$, so for the same flow rate the oil would absorb less than half the heat per degree of temperature rise. The engine would heat up further before the oil could remove the heat, and overheating would continue.

(b) One other property: the boiling point (water boils at 100°C, which sets the upper operating temperature; a coolant that boils too low would vapourise in the engine).

Other acceptable properties: freezing point, corrosion resistance to the engine metals, viscosity, non-flammability.

15. (a) The technician's advice is not physically sensible. The three thermometers disagree because they use three different thermometric substances (mercury expansion, platinum resistance, thermocouple EMF), and each has its own slight non-linearity with temperature. Recent calibration does not remove this inherent disagreement; it only fixes the readings at the two calibration points.

(b) The student should record the platinum resistance thermometer's reading of 48.5°C. Among the three instruments, the platinum thermometer has the smallest departure from true thermodynamic temperature between the fixed points, and its reading can be taken as the most accurate of the three.

16. When the sambusas are dropped into the oil, heat flows from the hot oil to the cold sambusas, cooling the oil. The oil's thermal capacity is $C = mc$, fixed by the mass of oil in the pot. The cooling caused by the sambusas is $\Delta T = Q/C$, where Q is the heat the sambusas extract. When twenty sambusas are added at once, Q is four times larger than when five are added, so the oil's temperature drops four times further. A larger drop means the sambusas cook in cooler oil and therefore take longer. The vendor can avoid the problem by using a larger pot of oil (higher thermal capacity) so that the same number of sambusas produces a smaller temperature drop, or by frying the sambusas in smaller batches so that each batch encounters oil close to the intended cooking temperature.

17. When the soda bottle is placed in the iced water basin, heat flows from the soda (and from the water) to the ice. As it receives heat, the ice begins to melt, absorbing the latent heat of fusion $L_f = 3.34 \times 10^5\,\mathrm{J\,kg^{-1}}$ without changing its own temperature. The ice therefore acts as an energy sink that keeps the water in the basin at 0°C until all of it has melted. In the 4°C basin, there is no latent-heat sink; every joule extracted from the soda raises the temperature of the surrounding water, reducing the temperature difference between the soda and the water and slowing the cooling rate. So the iced-water basin stays colder longer and cools the soda faster, just as Kipute claims.

18. When the lid is off, water molecules at the surface of the stew can escape as vapour by evaporation, and each escaping molecule carries away the latent heat of vapourisation. This heat is lost to the air and does not go into warming the stew. A significant fraction of the flame's heat is

consumed replacing the lost molecules' energy rather than raising the stew's temperature. With the lid on, the evaporated molecules accumulate in the air trapped above the stew and many of them condense back onto the surface, returning the heat. The lidded pot therefore reaches boiling sooner.

19. (a) The valve stays silent for the first ten minutes because the water inside the cooker is being warmed from room temperature toward its (elevated) boiling point, and no net steam is being produced yet. Once the water reaches its boiling point at the cooker's internal pressure, steam begins to accumulate, the internal pressure rises to the valve's set point, and the valve begins to hiss as it vents steam.

(b) The continued hissing means that the rate at which steam is being produced (by the flame's boiling of the water) exceeds the rate at which the valve can release it without the pressure rising further. The valve is therefore maintaining the pressure at its set value by matching its steam-release rate to the steam-production rate.

(c) If the flame is reduced but the valve continues to hiss, the cooker's internal temperature does **not** drop. The temperature is locked at the boiling point corresponding to the valve's set pressure, and so long as any steam is being produced and the valve is hissing, the pressure (and hence the temperature) stays at that set value. The reduction in flame strength reduces only the rate of steam production, and so the rate at which the valve hisses.

EXERCISE 2C

20. 0.47°C.

In a real waterfall, the actual temperature rise is smaller because: (1) some of the gravitational potential energy goes into kinetic energy of spray and air motion, which is dissipated to the surrounding atmosphere rather than into the falling water itself, and (2) some of the falling water evaporates on the way down, absorbing latent heat that would otherwise have raised the temperature.

21. 16.3°C

22. 21.2g, 100°C.

23. $a = 40\mu V/°C$, $b = -0.040\mu V/°C^2$, $1950\mu V$ or $1.95mV$

24. The rate of growth of an ice layer of thickness x is $dx/dt = k\Delta T/(\rho L_f x)$, derived in the chapter from balancing the rate of heat conducted through the ice with the rate at which the released heat of fusion is produced. Rearranging and integrating:

$$x\,dx = \frac{k\Delta T}{\rho L_f}\,dt$$

Integrating from $x_1 = 0.060m$ at time 0 to $x_2 = 0.070m$ at time t:

$$\int_{x_1}^{x_2} x\,dx = \frac{k\Delta T}{\rho L_f}\int_0^t dt$$

$$\frac{x_2^2 - x_1^2}{2} = \frac{k\Delta T}{\rho L_f}\,t$$

$$t = \frac{\rho L_f (x_2^2 - x_1^2)}{2k\Delta T}$$

Substituting the numerical values, with $\Delta T = 8K$:

$$t = \frac{920kg\,m^{-3} \times 3.34 \times 10^5 J\,kg^{-1} \times ((0.070m)^2 - (0.060m)^2)}{2 \times 2.2W\,m^{-1}K^{-1} \times 8K} = 1.135 \times 10^4 s \approx 3.15h$$

The ice layer grows from 6.0cm to 7.0cm in about 3.15h.

25. 46.4°C.

26. (a) 53.33°C (b) 52.45°C (c) 1.7%

27. 692g

28. (a) $1.68 \times 10^5 J$ (b) 1.4min or 84s

29. $0.80gh^{-1}$

30. Specific heat capacity of the solid phase (segment 1, $t = 0$ to 60s, $\Delta T = 40K$):

$$Q_1 = P \times t_1 = 50W \times 60s = 3000J$$

$$c_{solid} = \frac{Q_1}{m\,\Delta T_1} = \frac{3000J}{0.100kg \times 40K} = 750J\,kg^{-1}K^{-1}$$

Latent heat of fusion (segment 2, $t = 60s$ to 140s, $\Delta t_2 = 80s$):

$$Q_2 = P \times t_2 = 50W \times 80s = 4000J$$

$$L_f = \frac{Q_2}{m} = \frac{4000J}{0.100kg} = 4.0 \times 10^4 J\,kg^{-1}$$

Specific heat capacity of the liquid phase (segment 3, $t = 140s$ to 200s, $\Delta T = 70K$):

$$Q_3 = P \times t_3 = 50W \times 60s = 3000J$$

$$c_{liquid} = \frac{Q_3}{m\,\Delta T_3} = \frac{3000J}{0.100kg \times 70K} = 429J\,kg^{-1}K^{-1}$$

31. Kinetic energy of the bullet just before impact:

$$KE = 1/2\,mv^2 = 1/2 \times 0.005kg \times (350m\,s^{-1})^2 = 306.25J$$

Heat required to raise the bullet's temperature from 25°C to its melting point 327°C:

$$Q_{warm} = m\,c_{Pb}(327 - 25)K = 0.005kg \times 128J\,kg^{-1}K^{-1} \times 302K = 193.28J$$

Since $KE = 306.25J > Q_{warm} = 193.28J$, the bullet reaches its melting point and continues to absorb heat. The remaining energy goes into melting some of the lead at 327°C:

$$Q_{remaining} = KE - Q_{warm} = 306.25J - 193.28J = 112.97J$$

Mass of lead that melts:

$$m_{melt} = \frac{Q_{remaining}}{L_{f,Pb}} = \frac{112.97J}{2.5 \times 10^4 J\,kg^{-1}} = 4.52 \times 10^{-3}kg = 4.52g$$

About 4.5g of the bullet's mass melts on impact.

32. 0.147g, 0.3%.

EXERCISE 3A

1. (a) Metal has much higher thermal conductivity than wood. In the brief contact, the metal conducts heat away from the skin rapidly, drawing enough thermal energy to cause a burn. Wood conducts heat so slowly that very little energy reaches the skin in half a second, so the temperature of the skin barely changes.

(b) Not correct. Both handles are at the same temperature (100°C, the temperature of boiling water). The student is confusing temperature (which is the same) with the rate of heat transfer (which depends on thermal conductivity). The metal feels hotter because it delivers heat to the hand faster, not because it is at a higher temperature.

2. (a) In an open air gap, although air itself is a poor conductor, the air is free to move. Warm air near the hot surface rises and is replaced by cooler air, setting up convection currents that carry heat across the gap far more effectively than conduction alone. Felt prevents this air movement.

(b) Convection. The felt traps air in small pockets between its fibres, preventing bulk air movement. The open gap allows free convection.

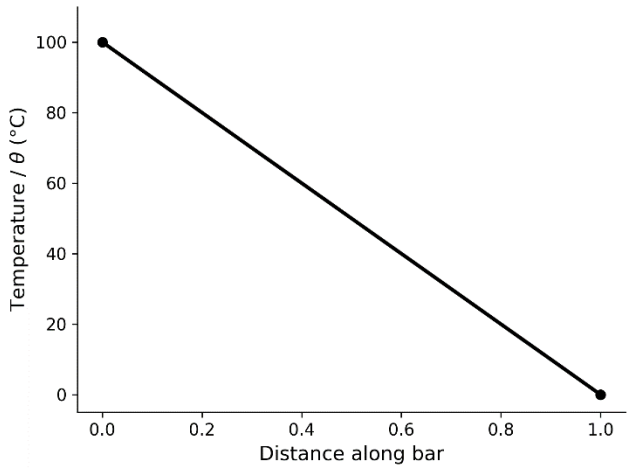
(c) All fibrous materials work by trapping air in small, isolated pockets. The fibres themselves conduct a small amount of heat, but their main function is to break the air into cells too small for convection currents to develop. The insulation comes from the still air, not from the fibre material itself.

3. (a) In the electrical analogy, temperature difference is analogous to voltage drop and heat flow rate is analogous to current. In a series circuit, the same current flows through every resistor, and the voltage drop across each is proportional to its resistance ($V = IR$). Similarly, the same heat flow rate passes through every layer, and the temperature drop across each layer is proportional to its thermal resistance ($\Delta\theta = \dot{Q} \times R$).

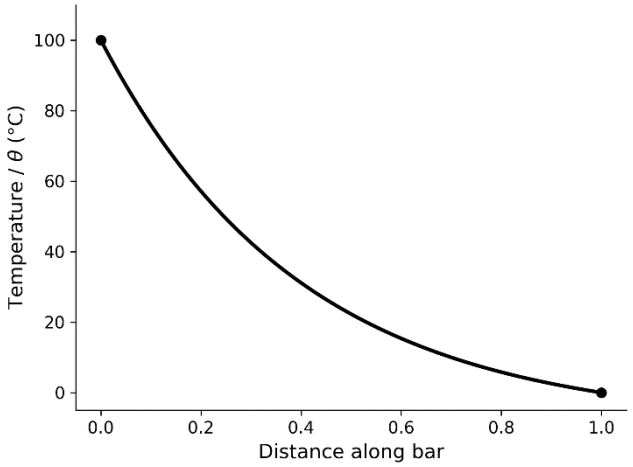
(b) The layer with the largest thermal resistance has the largest temperature drop. This is the layer with the combination of lowest thermal conductivity and greatest thickness, since $R = L/(kA)$ and A is constant for series arrangement. So it is not necessarily the thickest layer or the one with the lowest conductivity alone; it is the one with the highest L/k ratio.

(c) A thick layer of a good conductor (high k) can have a lower thermal resistance than a thin layer of a poor conductor (low k). Thickness matters, but only in combination with conductivity. The builder's claim ignores the role of k.

4. (a)



(b)



Explanation: When unlagged, heat leaks from the sides along the entire length. The heat flow rate decreases as you move from the hot end toward the cold end (because some heat has already escaped). A decreasing heat flow rate means a decreasing temperature gradient, so the profile curves.

Direction: It is concave upward (curves toward the hot end). This because the temperature drops steeply near the hot end (where the heat flow is large) and flattens near the cold end (where less heat remains).

(c) Below 50°C. The heat leaking from the sides throughout the bar means less heat reaches the midpoint than in the lagged case. With less heat flowing through the second half of the bar, the temperature at the midpoint is pulled downward.

5. (a) Newton's regime begins when the excess drops below about 30°C. Since the room is at 28°C, this occurs at about 58°C (excess = 58 − 28 = 30°C). Below 58°C, Newton's law applies; above 58°C, Dulong-Petit governs.

(b) Faster. The Dulong-Petit law gives a cooling rate that increases more rapidly than linearly with excess temperature (it follows a power law with index greater than 1). Newton's law, being linear in the excess, underestimates the cooling rate at large excess temperatures.

6. The student applied T^4 to the body's emission alone, ignoring the surroundings' contribution. The net heat loss is $\varepsilon\sigma A(T^4 - T_s{}^4)$, not $\varepsilon\sigma AT^4$. The ratio of net losses is not the same as the ratio of emitted powers.

7. (a) Convection. The hole allows warm air to rise out of the flask and be replaced by cooler air from outside, creating a convection current through the opening. This is the most immediate effect because convection can transport heat much faster than radiation or conduction through the stopper material.

(b) Radiation increases. The silvered surfaces had very low emissivity ($\varepsilon \approx 0.02$), suppressing radiative exchange across the vacuum gap. As the coating deteriorates, the emissivity rises, and the surfaces emit and absorb more radiation. Since the vacuum cannot carry heat by conduction or convection, radiation is the only mechanism operating across the gap, and its increase directly raises the overall heat-transfer rate.

(c) The student's error is assuming the flask has a preferred direction. The flask suppresses heat transfer in both directions by the same set of mechanisms. When hot coffee is inside, it slows heat flowing outward. When cold juice is inside, the same vacuum, silvering, and stopper slow heat flowing inward from the warm room. The flask reduces the rate of heat transfer regardless of direction.

8. (a) Block A cools faster initially. It has a larger excess temperature ($80 - 20 = 60$°C versus $60 - 20 = 40$°C), and Newton's law says the rate of temperature fall is proportional to the excess.

(b) Block A takes longer. It starts further from room temperature, so it has more excess to shed. Even though it cools faster at every moment (because it is always hotter than B at the corresponding elapsed time), the extra distance it must cover outweighs the faster speed.

(c) Yes, they have the same half-cooling time. The half-cooling time $t_{1/2} = \ln 2/\lambda$ depends only on λ, which is the same for both blocks (same mass, surface, material, and surroundings).

The answers to (a) and (b) are not contradictory: A cools faster in absolute terms (more degrees per minute) but not in relative terms (the same fraction of its excess decays in each half-cooling time). The exponential decay is the same; only the starting point differs.

9. (a) Prevost's theory: every body emits and absorbs radiation simultaneously. The baby's skin (at about 36°C) emits infrared radiation toward all surrounding surfaces. If the incubator walls are cooler than the baby's skin (for example, at 25°C), the walls emit less radiation back toward the baby than the baby emits toward them. The net radiative exchange is from the baby to the walls. The air temperature is irrelevant to radiation; radiation does not pass through air by convection but travels directly between surfaces.

(b) The problem is that the incubator walls are at a lower temperature than the baby. To eliminate the net radiative loss, the walls would need to be at the same temperature as the baby's skin (about 36°C). Then, by Prevost's theory, the baby would absorb exactly as much radiation from the walls as it emits, and the net radiative exchange would be zero.

(c) The double-walled design circulates warm water between the inner and outer walls, raising the inner wall temperature close to the baby's skin temperature. The baby now sees warm walls radiating strongly back toward it, reducing the net radiative loss to near zero. The air inside is also warm, so convective losses are minimal. Both mechanisms are addressed simultaneously.

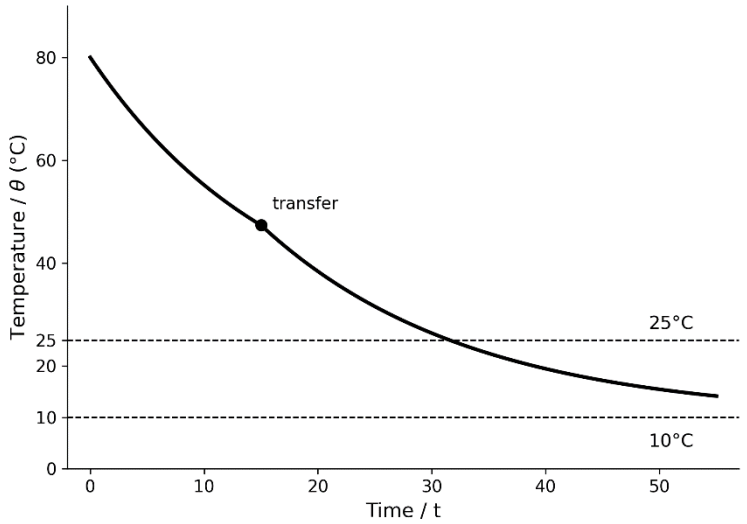
10. (a) Conduction. Heat flows from the liquid water at 0°C (just below the ice) upward through the solid ice to the cold air above. The water is the heat source; the air is the heat sink; the ice is the conducting path.

(b) Fourier's law: $\frac{dQ}{dt} = \frac{kA\Delta\theta}{x}$. The temperature difference $\Delta\theta$ and the area A stay the same, but the thickness x of the ice increases as more ice forms. A thicker ice layer means a longer conduction path, which increases the thermal resistance and decreases the heat flow rate.

(c) The heat that flows through the ice is the heat that freezes more water at the bottom surface. As the ice thickens, the heat flow rate drops (from (b)), so less water freezes per second, and the ice grows more slowly. The ice is insulating the water below from the cold air above. Each new layer of ice makes the insulation thicker, which slows the formation of the next layer. The process is self-limiting: the ice can never grow fast enough to overcome its own insulating effect.

11. (a) The rate of temperature fall increases instantly. In the 25°C room, the driving force was the excess over 25°C. In the 10°C room, the excess is suddenly larger (body temperature minus 10°C instead of minus 25°C), so Newton's law gives a larger rate of temperature fall at that instant.

(b) The cooling curve starts as a standard exponential decay toward 25°C. At the transfer point, the curve kinks downward: the slope becomes steeper and the curve now decays toward 10°C instead of 25°C.



(c) No. The cooling constant $\lambda = k/(mc)$ depends on k, which for radiative cooling is $k = 4\varepsilon\sigma AT_s{}^3$. Since T_s is different in the two rooms (298K versus 283K), the value of k and therefore λ is different. The cooling constant in the 10°C room is slightly smaller than in the 25°C room.

12. (a) Natural convection requires buoyancy: warm air must be less dense than cooler air and must rise under the influence of gravity. In microgravity (free fall around Earth), there is no effective gravitational force to drive density-driven flow. Warm air stays where it is and does not rise, so natural convection currents cannot form. Hot spots develop around equipment with no natural mechanism to carry the heat away.

(b) Forced convection. The fans mechanically push air over the hot surfaces, replacing warm air with cooler air regardless of gravity. Forced convection does not rely on buoyancy; it relies on an external force (the fan blades) to move the fluid. This works in any gravitational environment, including microgravity.

(c) Outside the station, there is a near-vacuum: no material medium exists for conduction (which requires molecular contact) or convection (which requires a fluid). Radiation is the only mechanism that can transport energy through a vacuum, because electromagnetic waves do not need a medium. The radiator panels are large, dark surfaces designed to maximise radiative emission into the cold of space (effective temperature about 3K), and they are the station's only route for rejecting waste heat.

EXERCISE 3B

13. (a) Marble has much higher thermal conductivity than wood. When the foot touches marble, heat flows rapidly from the warm skin into the cold marble, producing a strong sensation of cold. Wood conducts heat so slowly that very little energy leaves the foot in the same contact time, so the foot stays warm. The sensation of cold is caused by the rate of heat loss from the skin, not by the floor temperature.
(b) Rubber has low thermal conductivity (similar to wood). The sandal acts as an insulating layer between the foot and the marble, reducing the rate of heat flow from the skin. The foot loses heat much more slowly and feels less cold.
14. (a) Clay has low thermal conductivity compared to the high temperature gradient across it, but the fire is in direct contact with the clay body, supplying heat continuously. The wooden handle is separated from the fire and connected to the clay body through a narrow junction. Wood has very low thermal conductivity, so heat from the clay body conducts very slowly along the handle. The handle stays cool because the rate of heat arriving through the wood is too small to raise its temperature significantly.
(b) **First mechanism:** ***evaporation.*** The damp cloth loses water by evaporation, which draws latent heat from the handle and the cloth, cooling both. **Second mechanism**: ***conduction.*** The water in the cloth has higher thermal capacity than dry cloth, so it absorbs more heat before its temperature rises, acting as a temporary heat sink that delays the warming of the handle.
15. (a) No. The blanket generates no heat. It is a passive insulator that reduces the rate of heat transfer between the object it covers and the surroundings.
(b) **Conduction:** the blanket material and the air trapped within it have low thermal conductivity, resisting heat flow through the blanket's thickness. **Convection:** the blanket traps air in small pockets between its fibres, preventing convection currents from carrying heat across the layer. **Radiation:** the blanket's fabric interrupts the direct line of sight between the warm body (or ice) and the surroundings, blocking radiative exchange.
(c) Fourier's law is symmetric: the rate of heat flow depends on the temperature difference and the thermal resistance, regardless of which side is hotter. The blanket increases the thermal resistance between the object and the surroundings. This slows heat flow outward (keeping the person warm) and inward (keeping the ice cold) by the same mechanism. The blanket does not choose a direction; it resists heat flow in both.
16. (a) A black surface has high absorptivity (close to 1), so it absorbs most of the incoming solar radiation and converts it to thermal energy. A shiny aluminium surface reflects most of the sunlight and absorbs very little. Painting the panel black maximises the solar energy captured by the panel.
(b) Natural convection (thermosiphon effect). The sun heats the water in the panel, making it less dense. The warm water rises naturally through the pipes to the storage tank above. Cooler, denser water from the bottom of the tank flows down to the panel to replace it. This density-driven circulation continues as long as the panel is hotter than the tank, with no pump needed.
(c) On a clear day, the solar intensity reaching the panel is much higher than on an overcast day, because clouds reflect and scatter a large fraction of the incoming sunlight. The rate of energy absorption by the panel depends on the intensity of the radiation striking it, not on the air temperature. Even at the same air temperature, the panel receives less radiant energy on a cloudy day and therefore heats the water less.
17. (a) **Error 1:** the greenhouse effect does not prevent heat from escaping. It absorbs outgoing infrared and re-emits some of it back toward the surface, but the rest escapes to space. The atmosphere is not a sealed container. **Error 2:** the temperature does not rise indefinitely. The system reaches an equilibrium where the rate of outgoing radiation from the top of the atmosphere equals the rate of incoming solar absorption. At that equilibrium, the temperature stabilises.
(b) The surface warms until it radiates enough infrared energy that the total outgoing radiation from the top of the atmosphere balances the incoming solar absorption. Greenhouse gases shift this equilibrium to a higher surface temperature (because the surface must be warmer to drive enough radiation through the absorbing atmosphere), but the system always finds a new balance.
(c) Without greenhouse gases, the atmosphere would be transparent to all wavelengths. The surface would radiate directly to space with no atmospheric absorption or re-emission. The equilibrium temperature would drop to about 255K (−18°C), far below the freezing point of water. The planet's surface would be permanently frozen.
18. (a) The hot tea loses heat faster. Its excess temperature over the room ($80 - 20 = 60$°C) is much larger than the cold water's deficit ($20 - 10 = 10$°C). Newton's law of cooling says the rate of heat transfer is proportional to the temperature difference, so the tea's heat loss rate is about six times larger than the cold water's heat gain rate.
(b) The cold water gains heat. It is below room temperature, so heat flows from the warm room air into the cold water by convection and conduction, and from the warm surroundings by radiation.
(c) Both will eventually reach 20°C (room temperature). The cold water reaches it first because it has a much smaller temperature gap to close (only 10°C versus 60°C for the tea). Even though the tea cools faster in absolute terms (more watts of heat flow), it has a much longer journey to make.
19. (a) Conduction: add a layer of insulating material (such as foam board or ceiling board) beneath the iron roof. The insulation has very low thermal conductivity, creating a high thermal resistance between the hot roof and the interior, reducing the rate of heat conduction into the house.
(b) Convection: install a ventilated air gap (ridge vents at the top, eave vents at the bottom) between the roof and the ceiling. Hot air beneath the iron roof rises and escapes through the ridge vents, pulling cooler outside air in through the eaves. This removes heat by natural convection before it reaches the living space.
(c) Radiation: paint the outer surface of the iron roof with reflective white or silver paint, or install a reflective foil barrier beneath the roof. A reflective surface has low absorptivity, so it absorbs less solar radiation. A foil barrier beneath the roof has low emissivity, so it radiates less heat downward into the house.
20. (a) **Conduction:** newspaper has low thermal conductivity and the multiple layers trap air between them, increasing the total thermal resistance. **Convection:** the newspaper wrapping prevents warm air from rising freely from the pot surface and being replaced by cooler air, suppressing natural convection.
(b) Newton's law of cooling says the rate of temperature fall is proportional to the excess temperature over the surroundings. When the *ugali* is very hot, the excess is large and the rate of heat loss is high, so the temperature falls quickly. As the *ugali* cools, the excess shrinks, the driving force for heat loss weakens, and the rate of cooling slows. This self-slowing is the nature of exponential decay: the rate depends on the current excess, not on the initial temperature.
(c) Aluminium foil is a better reflector of infrared radiation (low emissivity), so it would reduce radiative heat loss more effectively than newspaper. However, aluminium is a good thermal conductor, so it would reduce conductive and convective heat loss less effectively than newspaper (which

traps insulating air layers). The best solution would be aluminium foil on the outside (to reflect radiation) with newspaper underneath (to trap air and resist conduction and convection).

21. (a) Yes, Kipute is correct that the tea in the metal mug cooled faster. Metal conducts heat rapidly from the tea through the mug wall to the outer surface, where it is lost to the air by convection and radiation. Ceramic conducts heat more slowly, so less heat reaches the outer surface per second.

(b) Thermal conductivity determines how fast heat passes through the wall, but it is not the only factor. The thermal capacity of the mug material also matters. A ceramic mug is typically thicker and has higher specific heat capacity than a thin metal mug, so the ceramic absorbs and stores more heat from the tea before its outer surface warms enough to lose heat rapidly. The metal mug, being thin with low thermal capacity, warms through almost instantly and begins losing heat immediately. The rate of cooling depends on both the conductivity (how fast heat reaches the surface) and the thermal mass of the mug (how much heat the mug itself absorbs along the way).

(c) Yes. If no heat can escape from either mug, both systems (tea + mug) would reach thermal equilibrium internally. The tea would cool slightly as it shares heat with the mug material, and the mug would warm. The final temperature in each case would depend on the relative thermal capacities of the tea and the mug, which are different for metal and ceramic. So the final temperatures would be slightly different, but both would be close to the initial tea temperature because the tea's thermal capacity is much larger than either mug's.

22. (a) Foam has very low thermal conductivity compared to concrete. Adding the foam layer places a high thermal resistance in series with the concrete. By Fourier's law, the rate of heat flow through the composite is determined by the total thermal resistance, which is now much larger than concrete alone. The foam reduces the rate of heat conduction through the roof dramatically.

(b) In a series arrangement, the temperature drop across each layer is proportional to its thermal resistance. Foam has much higher thermal resistance per unit thickness than concrete (because its conductivity is much lower). Even though the concrete is thicker, its low resistance means only a small temperature drop occurs across it. Most of the total temperature difference between outside and inside appears across the thin foam layer, because that is where the resistance is concentrated.

(c) The foam's thermal resistance works identically regardless of which side is hotter. During the day, it resists heat flowing inward from the hot roof to the cool interior. At night, it resists heat flowing outward from the warm interior to the cold roof. The resistance is a property of the material, not of the direction of heat flow. This is a direct consequence of Fourier's law being symmetric with respect to the sign of the temperature difference.

EXERCISE 3C

23. $76\text{Wm}^{-1}\text{K}^{-1}$

24. 12.1Wm^{-2}, $18.7°\text{C}$, 4.6

25. From Newton's law of cooling, the time for each interval:

$$t_1 = \frac{1}{\lambda}\ln\left(\frac{\theta_1 - \theta_s}{\theta_2 - \theta_s}\right)$$
$$t_2 = \frac{1}{\lambda}\ln\left(\frac{\theta_2 - \theta_s}{\theta_3 - \theta_s}\right)$$

Since the drops are equal, $\theta_1 - \theta_2 = \theta_2 - \theta_3 = d$ (say). Then $\theta_3 = \theta_2 - d$ and $\theta_1 = \theta_2 + d$. The excess temperatures are:
$\theta_1 - \theta_s = \theta_2 - \theta_s + d$ and $\theta_3 - \theta_s = \theta_2 - \theta_s - d$
The argument of the logarithm in t_2 is:

$$\frac{\theta_2 - \theta_s}{\theta_2 - \theta_s - d}$$

and in t_1:

$$\frac{\theta_2 - \theta_s + d}{\theta_2 - \theta_s}$$

Let $E = \theta_2 - \theta_s$ (the excess at the midpoint). Then:

$$\frac{t_2}{t_1} = \frac{\ln\left(\frac{E}{E-d}\right)}{\ln\left(\frac{E+d}{E}\right)}$$

Since $E > d$ (the body is above the surroundings throughout), we have $\frac{E}{E-d} > \frac{E+d}{E}$ because subtracting d from a smaller denominator produces a larger ratio than adding d to a larger numerator. Therefore $\ln\left(\frac{E}{E-d}\right) > \ln\left(\frac{E+d}{E}\right)$, giving $t_2 > t_1$.

26. (a) The exact equation from Stefan-Boltzmann:

$$mc\frac{dT}{dt} = -\varepsilon\sigma A\left(T^4 - T_s{}^4\right)$$
$$\frac{dT}{dt} = -\frac{\varepsilon\sigma A}{mc}\left(T^4 - T_s{}^4\right)$$

(b) The right-hand side contains T^4, not $(T - T_s)$. The equation is not of the form $dT/dt = -\lambda(T - T_s)$, so the solution is not a simple exponential. The T^4 dependence makes the equation nonlinear: the rate of cooling changes disproportionately as T changes, unlike Newton's law where the rate is strictly proportional to the excess.

(c) 0.046Ks^{-1}

27. (a) 139W (b) 2.7W

28. $6.54 \times 10^{-8}\ \text{Wm}^{-2}\text{K}^{-4}$

29. 1.81kW, 1450nm

30. (a) Furnace: 24200K, Star: 6042K (b) $P_{\text{furnace}}: P_{\text{star}} = 257:1$ (c) 9350nm (or 9.35μm).

31. (a) Rate of heat loss per unit area through ice of thickness x:

$$\frac{dQ}{dt} = \frac{kA\Delta T}{x}$$

This heat freezes a thin layer dx of water:

$$dQ = \rho_{\text{ice}}\, A\, dx \times L_f$$

So:

$$\rho_{ice}\, A\, L_f \frac{dx}{dt} = \frac{kA\Delta T}{x}$$

$$x\, dx = \frac{k\Delta T}{\rho_{ice} L_f}\, dt$$

Integrating from x_1 to x_2:

$$\int_{x_1}^{x_2} x\, dx = \frac{k\Delta T}{\rho_{ice} L_f} \int_0^t dt$$

$$\frac{1}{2}({x_2}^2 - {x_1}^2) = \frac{k\Delta T}{\rho_{ice} L_f}\, t$$

$$t = \frac{\rho_{ice} L_f}{2k\Delta T}({x_2}^2 - {x_1}^2)$$

(b) Time ratio:

$$\frac{t_{10\to11}}{t_{1\to2}} = \frac{(0.11)^2 - (0.10)^2}{(0.02)^2 - (0.01)^2} = 7.0$$

Growing from 10cm to 11cm takes 7 times longer than growing from 1cm to 2cm.

32. Two cylindrical shells in series. Radial thermal resistance of each:

$$R = \frac{\ln\left(\frac{r_2}{r_1}\right)}{2\pi kL}$$

Copper shell ($r_1 = 0.020$m, $r_2 = 0.030$m):

$$R_{Cu} = \frac{\ln(0.030/0.020)}{2\pi \times 380 \times 5.0} = \frac{0.4055}{11938} = 3.40 \times 10^{-5}\ KW^{-1}$$

Asbestos shell ($r_1 = 0.030$m, $r_2 = 0.070$m):

$$R_{asb} = \frac{\ln(0.070/0.030)}{2\pi \times 0.08 \times 5.0} = \frac{0.8473}{2.513} = 0.337 KW^{-1}$$

$$R_{total} = 3.40 \times 10^{-5} + 0.337 \approx 0.337 KW^{-1}$$

(Copper resistance is negligible.)

$$\frac{dQ}{dt} = \frac{58K}{0.337 KW^{-1}} = 172W$$

Temperature at copper-asbestos boundary:

$$\theta_{boundary} = 80°C - 172W \times 3.40 \times 10^{-5}\ KW^{-1} = 80°C - 0.006°C \approx 80.0°C$$

(The copper wall is essentially isothermal: the entire temperature drop occurs across the asbestos.)

33. (a) With albedo, the power absorbed by Earth is:

$$P_{abs} = (1 - \alpha)\varphi_s\, \pi {R_e}^2$$

The power radiated by Earth as a blackbody sphere:

$$P_{rad} = \sigma(4\pi {R_e}^2){T_e}^4$$

At equilibrium:

$$(1 - \alpha)\varphi_s\, \pi {R_e}^2 = \sigma \times 4\pi {R_e}^2 {T_e}^4$$

$$\boldsymbol{{T_e}^4 = \frac{(1 - \alpha)\varphi_s}{4\sigma}}$$

With $\alpha = 0$ (no reflection), this gives the bare-blackbody value. With $\alpha = 0.30$, the factor $(1 - \alpha) = 0.70$ reduces the numerator, so ${T_e}^4$ is smaller and therefore T_e is lower than the bare-blackbody value.

(b) Ice and snow have high albedo (0.8 to 0.9); ocean water has low albedo (0.06 to 0.10). If ice caps melt, they are replaced by dark ocean, which absorbs more sunlight. The overall planetary albedo decreases. With lower α, the factor (1 − α) increases, so more solar energy is absorbed, and the equilibrium temperature rises further.

(c) This is a destabilising (positive) feedback. A small initial warming melts some ice, which lowers the albedo, which causes more warming, which melts more ice, which lowers the albedo further. Each step amplifies the initial perturbation rather than counteracting it. This is called the ice-albedo feedback and is one of the most important positive feedbacks in climate physics.

EXERCISE 4A

1. Errors:

First error: The assumption that molecules touch each other at all; in the ideal gas model the molecules are taken as essentially point masses that interact only briefly at collisions.

Second error: Identifying the cause of pressure with collisions between molecules, whereas pressure is produced by collisions of the molecules with the walls of the container.

Correct reasoning:

Compressing the gas reduces the container volume, so each molecule travels a shorter distance between successive wall collisions and strikes the walls more frequently per second; the increased rate of momentum transfer to the walls is the true cause of the pressure increase.

2. (a) Intermolecular forces are assumed to be negligible except during the instant of a collision, so there is no force acting on a molecule in the intervals between collisions. With no force there is no acceleration, and by Newton's first law the molecule moves in a straight line at constant velocity.

(b) Assuming the collision time is negligible allows us to treat each molecule as moving freely between instantaneous collisions. The assumption simplifies the derivation of pressure by making it unnecessary to track the detailed mechanics of what happens during a collision and enables the momentum change at the wall to be treated as a sudden reversal.

3. The average rate at which molecules strike a square centimetre of wall is the same in both containers. Pressure is force per unit area on the wall, and the two containers have the same pressure; since all walls experience the same rate of momentum transfer per unit area, the rate at which molecules strike one square centimetre is the same on any wall, whether flat or curved, in either container.

4. (a) The average kinetic energy per molecule of an ideal gas is proportional to the Kelvin temperature, not the Celsius temperature, and the two scales have different zeros. Doubling the Celsius temperature does not double the Kelvin temperature, so the kinetic energy does not double.

(b) The student's rule requires two conditions simultaneously:

Doubling in Celsius ($\theta_2 = 2\theta_1$ in °C) and

Doubling in Kelvin ($\theta_2 + 273 = 2(\theta_1 + 273)$ in K).

Substituting the first into the second gives $2\theta_1 + 273 = 2\theta_1 + 546$, which reduces to 273 = 546, a contradiction. No pair of Celsius temperatures satisfies both conditions, so the student's rule never gives the correct ratio of kinetic energies exactly.

The rule becomes approximately correct in the limit of very high temperatures, where $\theta_1 \gg 273$°C. In this limit the additive offset of 273 is negligible compared with the Celsius values themselves, so the Celsius and Kelvin temperatures are approximately equal and their ratios are approximately the same. At stellar temperatures (thousands of kelvin) the rule holds to good accuracy.

5. (a) The average kinetic energy per molecule and the root mean square speed are the same in both containers, because both depend only on the Kelvin temperature.

(b) The pressure and the total internal energy are both twice as large in container B, because both scale in proportion to the number of moles (and hence the number of molecules) at the same temperature.

(c) An intensive quantity, such as temperature, pressure, or kinetic energy per molecule, is independent of the amount of matter present. An extensive quantity, such as total internal energy or total number of molecules, scales in proportion to the amount of matter. Doubling the amount of gas doubles the extensive quantities and leaves the intensive quantities unchanged.

6. The student has confused two different statistical operations. The root mean square speed is obtained by first squaring each speed, then averaging the squares, then taking the square root; the simple arithmetic mean averages the speeds directly without squaring. Whenever the molecular speeds differ from one another, squaring gives extra weight to the larger speeds, and the rms speed exceeds the arithmetic mean. The two would agree only in the special case where every molecule in the sample moves at exactly the same speed.

7. (a) A diatomic molecule consists of two atoms connected by a bond, and we can take the line joining the atoms as the bond axis. Rotation about an axis perpendicular to the bond carries the atoms in circular paths of finite radius and stores energy. Rotation about the bond axis itself, however, rotates the atoms around their own centres and stores negligible energy because the moment of inertia about that axis is very small (the atoms lie essentially on the axis). Only two of the three possible rotational axes therefore store appreciable kinetic energy, giving two rotational degrees of freedom.

(b) A polyatomic molecule such as water vapour has its atoms arranged in a non-linear configuration, so no one axis passes through every atom. All three perpendicular rotational axes carry the atoms in circular paths of non-zero radius and have non-zero moments of inertia. All three rotations therefore store appreciable kinetic energy, giving three rotational degrees of freedom.

8. The student's error is to identify work done on the gas with a rise in internal energy. The internal energy of an ideal gas depends only on its absolute temperature, not on its volume or pressure. Since the temperature is held constant during the compression, the internal energy does not change; the work done on the gas does not accumulate as increased molecular kinetic energy, because molecular kinetic energy is fixed by the temperature alone.

9. (a) For a real gas, the pressure rises by more than a factor of ten because the volume available to the molecules shrinks faster than the container volume: the volume occupied by the molecules themselves, negligible at the original volume, becomes a significant fraction of the container at the smaller volume. The assumption of negligible molecular volume (assumption 2 of the kinetic theory) fails, so the real gas pressure exceeds what the ideal gas equation predicts.

(b) When the compressed gas is cooled to a very low temperature, the attractive intermolecular forces begin to dominate over the slower-moving molecules, and the gas may liquefy. The ideal gas model, which assumes zero intermolecular forces, cannot describe liquefaction because in the absence of attractive forces there is nothing to hold molecules together into a liquid phase.

10. The rms speed of an ideal gas is proportional to the square root of the Kelvin temperature:

$$v_{rms} \propto \sqrt{T}$$

The factor by which it changes is therefore:

$$\frac{v_{rms_2}}{v_{rms_1}} = \sqrt{\frac{T_2}{T_1}} = \sqrt{\frac{1200\text{K}}{300\text{K}}} = 2$$

The root mean square speed doubles.

EXERCISE 4B

11. (a) At the higher temperature the nitrogen or LPG molecules inside the cylinder move with a greater mean square speed. Each molecule delivers a larger momentum change when it strikes a wall, and each molecule strikes the walls more frequently per second. Both effects multiply the rate of momentum transfer to the walls, and the pressure, which is this rate per unit area, rises in direct proportion to the Kelvin temperature at fixed volume.

(b) A cylinder stored in a closed, hot room has no way to lose heat to the surroundings and its internal temperature can climb further; the pressure would rise with it, possibly exceeding the safety rating and causing the cylinder to rupture. A shaded, well-ventilated place keeps the cylinder near ambient temperature by both blocking solar radiation and allowing convective heat loss to circulating air, keeping the internal pressure safely below the rating.

12. What has changed with altitude is the total atmospheric pressure, which at the summit of Kilimanjaro is about half its sea-level value. Since the fractional composition of air has not changed, the partial pressure of oxygen is also about half its sea-level value. The rate at which oxygen molecules diffuse from the lung air into the blood is driven by the partial-pressure gradient between the two, and with that gradient halved, the rate of oxygen uptake is reduced correspondingly. The mountain climber must breathe harder and more frequently to move the same amount of oxygen into his blood, and at the summit the body's demand begins to outstrip what his breathing can deliver.

13. (a) As the contents of the cooker are heated, the water molecules in the vapour phase move with greater mean square speed and strike the walls of the cooker more frequently and with greater momentum change per collision. The pressure of the vapour, which is the rate of momentum transfer to the walls per unit area, rises in direct proportion to the Kelvin temperature of the vapour at fixed volume.

(b) The pressure reaches a stable value when the safety valve opens and begins releasing steam. The escape of molecules through the valve removes them from the container and prevents the number density (and hence the pressure) from rising further. The valve maintains a near-constant pressure by venting just enough steam to balance the rate at which more water vapourises from the heated contents.

14. (a) Dalton's law states that each gas in a mixture exerts the same partial pressure as it would if it alone occupied the container at the same temperature. At sea level the total pressure is one atmosphere and the partial pressure of nitrogen is about 78% of this, roughly 0.78 atmospheres. At 20 metres depth the diver breathes air compressed to three atmospheres; the fractional composition has not changed, so each species' partial pressure is multiplied by three, giving a nitrogen partial pressure of about 2.34 atmospheres, approximately three times its surface value.

(b) At higher partial pressure, the number density of nitrogen molecules in the blood and tissues is correspondingly higher, so the rate at which nitrogen molecules collide with receptors, cell membranes, and other biological targets is also correspondingly higher. Even if each individual collision has negligible chemical effect, the cumulative rate of such collisions becomes large enough to disturb nerve transmission and produce the narcotic symptoms observed. Physiological activity of an inert gas depends on collision rates, not on chemistry.

15. Bakari is correct that air is a mixture of several gases, mainly nitrogen and oxygen, not a single kind of atom; Ally's wording is sloppy. But Ally's physical intuition is correct. At the same temperature, helium atoms (molar mass 4 g/mol) and the average air molecule (effective molar mass about 29 g/mol) carry the same average kinetic energy per molecule. A container filled with helium at a given temperature and pressure therefore holds the same number of molecules per unit volume as a container of air at the same temperature and pressure, but each helium atom is about seven times lighter than the average air molecule. The helium-filled balloon displaces a volume of air whose mass greatly exceeds the mass of the helium inside, and the buoyant force exceeds the weight of the balloon, causing it to rise. Ally's statement captures this mass-per-molecule comparison, even if the language of "atoms of air" is imprecise.

16. (a) Argon is monatomic, with only three translational degrees of freedom, while air consists mainly of diatomic nitrogen and oxygen molecules with five degrees of freedom (three translational plus two rotational). An argon atom therefore carries less thermal energy per collision than an air molecule and transports heat across the gap less effectively. In addition, argon atoms are heavier (molar mass 40 g/mol) than the average air molecule (29 g/mol), so at the same temperature the argon atoms move more slowly and arrive at the opposite pane less frequently. Both effects reduce the thermal conductivity of argon below that of air, making argon a better insulator.

(b) Krypton is also monatomic, so its number of degrees of freedom is the same as argon's (three). The insulation would not improve through that property. However, krypton atoms are much heavier than argon atoms (molar mass 84 g/mol compared with 40 g/mol), so krypton atoms move even more slowly at the same temperature and transport heat across the gap even less rapidly. The mass effect improves, the degrees-of-freedom effect does not; on balance, krypton provides better insulation than argon, which is why premium high-efficiency windows use it.

17. When Kipute pushes the handle down, the piston reduces the volume of air trapped in the pump cylinder rapidly. Because the compression is fast, the air has no time to lose heat to the surroundings through the walls; the compression is effectively an insulated one. The work done on the air during compression goes entirely into raising the kinetic energy of the air molecules, which means the temperature and the mean square speed of the molecules inside the pump have increased. These faster, hotter molecules then transfer their extra thermal energy by direct contact to the inner wall of the pump barrel, warming the metal. The metal feels warm not because of any external source of heat but because Kipute's muscular work has been converted, through the compression, into thermal energy of the gas and thence of the pump.

18. $T_{well} = 24 + 273 = 297\text{K}$; $T_{street} = 36 + 273 = 309\text{K}$.

Since $v_{rms} \propto \sqrt{T}$ for a given gas, the ratio of rms speeds is:

$$\frac{v_{rms_{street}}}{v_{rms_{well}}} = \sqrt{\frac{T_{street}}{T_{well}}} = \sqrt{\frac{309\text{K}}{297\text{K}}} = 1.020$$

The molecules on the hot street move only about 2% faster on average than those in the cool well. The word "much" is **inappropriate**: a 12°C difference in Celsius looks substantial in human terms, but on the Kelvin scale it represents only a 4% change in absolute temperature, and molecular speeds respond to the square root of that, giving barely 2%. The student has fallen into the familiar trap of treating Celsius differences as though they were fractional changes in the physically meaningful quantity, which is the Kelvin temperature.

EXERCISE 4C

19. Find the amount of nitrogen present from PV = nRT:

$$n = \frac{P_1 V}{RT_1} = \frac{(8.0 \times 10^6\text{Pa}) \times (0.040\text{m}^3)}{(8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (290\text{K})} = 132.7\text{mol}$$

At constant volume the heat supplied equals the change in internal energy. Nitrogen is diatomic (f = 5), so $\Delta U = (f/2)nR\Delta T$:

$$Q = \frac{5}{2} \times (132.7\text{mol}) \times (8.314\text{Jmol}^{-1}\text{K}^{-1}) \times (110\text{K}) = 3.03 \times 10^5\text{J}$$

The heat supplied to the gas is approximately 3.03×10^5J.

20. Equating the right-hand sides of $PV = \frac{1}{3}Nm\bar{c}^2$ and $PV = NkT$:

$$\frac{1}{3}Nm\bar{c}^2 = NkT$$

$$\bar{c}^2 = \frac{3kT}{m}$$

Taking the square root and using $m = M/N_A$ together with $R = k \times N_A$:

$$v_{rms} = \sqrt{\frac{3kT}{m}} = \sqrt{\frac{3RT}{M}}$$

For two gases A and B at the same temperature T, dividing one expression by the other gives:

$$\frac{v_{rms_A}}{v_{rms_B}} = \sqrt{\frac{M_B}{M_A}}$$

The ratio of rms speeds therefore depends only on the molar masses, not on temperature: raising the temperature accelerates both gases equally in the sense that their kinetic energies remain equal, so their speeds remain in the same mass-determined ratio.

21. The neon cylinder by a factor of 1.05.

22. The gas is helium. The identification is unique among common gases because helium is the only gas at A-Level with a molar mass of 4.0g/mol: the next heaviest common gas is neon at 20g/mol, five times heavier, far outside any reasonable rounding of the calculated value.

23. Newton's law of cooling, integrated, gives:

$$T(t) = T_s + (T_1 - T_s)\,e^{-\lambda t}$$

With $T_1 = 400\text{K}$, $T_s = 300\text{K}$, $\lambda = 2.0 \times 10^{-3}\text{s}^{-1}$, and $t = 600\text{s}$: $\lambda t = 1.2$ and $e^{-1.2} = 0.301$:

$$T(600s) = 300K + (100K) \times 0.301 = 330.1K$$

At constant volume and fixed n, the pressure is proportional to the Kelvin temperature:

$$P_2 = P_1 \times \frac{T_2}{T_1} = (5.0 \times 10^5 Pa) \times \frac{330.1K}{400K} = 4.13 \times 10^5 Pa$$

The pressure inside the container 10 minutes after placement is approximately 4.13×10^5 Pa.

24. Initial total moles: n = 1.5 + 2.0 + 0.5 = 4.0mol. Final total moles (after removal of water vapour): n' = 3.5mol. Pressure at constant V and T depends only on total moles:

$$\frac{\Delta P}{P} = \frac{n, - n}{n} = \frac{3.5 - 4.0}{4.0} = -0.125 = -12.5\%$$

The pressure decreases by 12.5%. Internal energy before removal, using f = 3 for He, f = 5 for N_2, and f = 6 for water vapour:

$$U_i = \frac{RT}{2} \times (3 \times 1.5mol + 5 \times 2.0mol + 6 \times 0.5mol) = \frac{RT}{2} \times 17.5mol$$

$$U_f = \frac{RT}{2} \times (3 \times 1.5mol + 5 \times 2.0mol) = \frac{RT}{2} \times 14.5mol$$

Percentage change in internal energy:

$$\frac{\Delta U}{U} = \frac{14.5 - 17.5}{17.5} = -0.1714 = -17.1\%$$

The internal energy decreases by 17.1%, considerably more than the 12.5% pressure drop. The two percentages differ because pressure at fixed V and T depends only on the total moles, so removing any species reduces the pressure in proportion to the moles removed. Internal energy, however, counts each species weighted by its degrees of freedom; water vapour is polyatomic (f = 6) and carries more energy per mole than either helium or nitrogen, so removing it takes away a larger fraction of the total internal energy than of the total moles.

25. From P = αV and PV = nRT, substituting P = αV:

$$\alpha V \times V = nRT, \text{ so } T = \frac{\alpha V^2}{nR}$$

The absolute temperature is therefore proportional to V^2. When the volume doubles, the temperature becomes $(2V_0/V_0)^2 = 4$ times its original value:

$$\frac{T_f}{T_0} = (2)^2 = 4$$

That is, $T_f = 4T_0$. The rms speed scales as the square root of the Kelvin temperature:

$$\frac{v_{rms_f}}{v_{rms_0}} = \sqrt{\frac{T_f}{T_0}} = \sqrt{4} = 2$$

The rms speed has doubled.

For a monatomic gas U = (3/2)nRT, so U is directly proportional to T:

$$\frac{U_f}{U_0} = \frac{T_f}{T_0} = 4$$

The final internal energy is four times the initial internal energy.

EXERCISE 5A

1. Both forms of the first law are correct statements of energy conservation; they differ only in the sign convention adopted for W. In the form $\Delta U = Q - W$, the symbol W denotes the work done by the gas on the surroundings: positive when the gas expands, negative when it is compressed. In the form $\Delta U = Q + W$, the symbol W denotes the work done on the gas by the surroundings: positive when the gas is compressed, negative when it expands. The two values of W differ only by a sign, $W_{by} = -W_{on}$, so the two equations are algebraically identical: heat in minus work-done-by equals heat in plus work-done-on, and both equal ΔU. The students are therefore both correct, provided each is consistent within their chosen convention; the disagreement is bookkeeping, not physics.

2. The error is to assume that heat exiting a gas must reduce its internal energy. For an ideal gas the internal energy depends only on temperature, $U = U(T)$, so as long as T is held constant, $\Delta U = 0$ regardless of how much heat flows in or out. During the isothermal compression the surroundings do positive work on the gas, raising the internal energy, while heat simultaneously leaves the gas, lowering it; with $\Delta U = 0$, the two effects must cancel exactly. The first law $\Delta U = Q - W$ (with W done **by** the gas, hence negative in compression) gives $Q = W < 0$, the heat removed equals in magnitude the work done on the gas. The work added to the gas leaves the gas as heat, not as a rise in internal energy.

3. The internal energy U is a state function: its value at any equilibrium state is determined by the state alone (for an ideal gas, by the temperature alone), independently of how the gas reached that state. At state A the gas has a specific internal energy U_A; at state B, a specific U_B; the change $\Delta U = U_B - U_A$ is fixed by the two endpoints and is independent of the path joining them. The work done by the gas, $W = \int PdV$, is the area under the chosen path on the PV diagram; different paths enclose different areas and therefore yield different values of W. The heat exchanged then follows from the first law as $Q = \Delta U + W$, and since ΔU is path-independent while W is not, Q also depends on the path. Both W and Q are path functions; only ΔU is a state function.

4. In a gas, heating at constant pressure produces a substantial volume change, since gases expand strongly with temperature. Part of the heat input is therefore expended as work done by the gas on the surroundings, $W = P\Delta V$, on top of raising the internal energy. At constant volume no such work is done, so all of the heat goes into raising U. The difference between the two heat capacities, $C_p - C_v = R$ per mole for an ideal gas, accounts exactly for the PΔV term. In a solid or liquid the thermal expansion coefficient is much smaller (of order $10^{-4}K^{-1}$ for liquids, $10^{-5}K^{-1}$ for solids), so the volume change associated with a temperature rise is small, and the work done against atmospheric pressure during heating is negligible compared with the heat absorbed. The difference $C_p - C_v$ does exist in principle for solids and liquids but is too small to matter for practical calculations, so a single specific heat capacity is quoted for each substance. The student's observation is therefore explained by the magnitude of thermal expansion, not by any difference in the underlying physics.

5. For any ideal gas the molar heat capacity at constant volume is $C_v = \frac{f}{2}R$, where f is the number of accessible degrees of freedom. Equal C_v therefore implies equal f. So far the student's chain of inference is sound: A and B share the same number of energy-storage modes per molecule. What this conclusion does **not** establish is the molecular identity of the gases, since many different molecules share the same f. Two distinct diatomic gases at room temperature, for example, both have $f = 5$ (three translational + two rotational) and therefore the same C_v: nitrogen and

oxygen are both diatomic but are chemically distinct; the helium isotopes ^{3}He and ^{4}He are both monatomic with $f = 3$ but differ only in their nuclear mass. To distinguish A from B beyond their degree-of-freedom count, a further measurement is required: the molar mass (which follows from the density at known temperature and pressure via $\rho = PM/RT$), the absorption spectrum (which is unique to each molecule), or a chemical assay. The molar C_v tells you how many ways a molecule stores energy, not which molecule it is.

6. The vessel is welded, so its internal volume is fixed; the gas it contains undergoes an isochoric cooling. For a fixed mass of ideal gas at constant volume the equation of state $PV = nRT$ reduces to $P/T = \text{constant}$. Converting the temperatures to the Kelvin scale gives $T_1 = 305\text{K}$ and $T_2 = 295\text{K}$, hence:

$$\frac{P_2}{P_1} = \frac{T_2}{T_1} = \frac{295\text{K}}{305\text{K}} = 1 - \frac{10}{305}$$

The pressure falls by exactly $10/305$ of its original value, neither more nor less. The assumption that underlies this exact result is the **ideal-gas** behaviour of the contents of the vessel over the modest 10K temperature interval. For typical industrial gases at temperatures well above their boiling points and at moderate pressures, this assumption is well justified.

7. By Avogadro's law, which follows from the ideal-gas equation of state $PV = nRT$, equal volumes of any gas at the same temperature and pressure contain the same number of moles, and therefore the same number of molecules. Both balloons have the same volume, the same temperature, and the same pressure, so both contain the same number N of molecules. The mass difference between the two balloons originates entirely in the molar mass of the gas filling them: helium is monatomic with molar mass about 4gmol^{-1}, while air is a mixture dominated by diatomic nitrogen (28gmol^{-1}) and oxygen (32gmol^{-1}), with an effective molar mass of about 29gmol^{-1}. The number of molecules per unit volume is the same; the mass of each molecule is not. The molecular property responsible for the lighter helium balloon is therefore the molar mass, not any difference in molecular count.

8. The error is to assume that any heat flow must change the temperature. Heat input raises the temperature only if the gas's internal energy is forced to rise; when the gas can simultaneously do work on its surroundings, the heat input emerges as work output rather than as a temperature rise. During an isothermal expansion of an ideal gas, $\Delta U = 0$ (since U depends only on T), and the first law reduces to $Q = W$: every joule of heat absorbed is converted to a joule of work done by the gas on the piston. Heat does flow, in substantial amounts; it just leaves the gas as mechanical work rather than appearing as kinetic energy of the molecules. In the isothermal case the heat flow keeps the molecules at constant average kinetic energy while the gas expands, the heat input being passed straight through to the surroundings as work.

9. The claim is partially correct. A genuinely fast process does prevent heat from flowing through the container walls during the brief duration of the change, so to a good approximation $Q \approx 0$; this is what 'adiabatic' means in the heat-isolation sense. The relation $PV^\gamma = \text{constant}$, however, requires more. Its derivation assumes that at every instant during the process the gas is in thermodynamic equilibrium, so that a single pressure and temperature can be defined for the whole sample, and the work done in an infinitesimal step is $dW = PdV$. This is the quasi-static condition. A truly rapid event, such as a tyre suddenly bursting or a piston released from rest, can be heat-isolated ($Q \approx 0$) but is not quasi-static; the gas is far from equilibrium during the event, P and T vary across the sample, and the relation $PV^\gamma = \text{constant}$ does not apply. The condition that must be added to 'happens rapidly' is therefore that the process be quasi-static: slow enough on the equilibration timescale of the gas, but fast enough on the heat-conduction timescale of the walls.

10. Both compressions begin at the same initial point A on the PV diagram and end at the same final volume $V_2 < V_1$. The isothermal follows the curve $PV = \text{constant}$, a hyperbola, and ends at point B with pressure $P_B = P_A V_1/V_2$. The adiabat follows $PV^\gamma = \text{constant}$, a steeper hyperbola, and ends at point C with pressure $P_C = P_A(V_1/V_2)^\gamma$. Since $\gamma > 1$, $P_C > P_B$: the adiabatic compression reaches the higher final pressure.

Sketch description (the student's hand-drawn version): an axis labelled V horizontally, P vertically; a common starting point A; two downward curves drawn to the left, both ending at $V = V_2$; the upper curve is the adiabat (steeper, ending higher), the lower curve is the isotherm (shallower, ending lower); end-states marked B (isothermal) and C (adiabatic) with $P_C > P_B$.

The molecular reason: in the isothermal compression heat must leave the gas to keep T constant, so the molecular kinetic energy stays unchanged and the pressure rises only as much as the volume reduction itself compels through $PV = \text{constant}$. In the adiabatic compression no heat leaves, so the work done on the gas raises both the internal energy and the temperature; the higher temperature gives faster, more frequent wall collisions, and the pressure rises more steeply than the volume reduction alone would predict.

11. (a) The net work done by the gas over a complete cycle is $W_{net} = \oint PdV$ ($\oint$ means integration over a complete closed cycle), which equals the area enclosed by the cycle on the PV diagram with sign. By the standard sign convention (clockwise positive), counter-clockwise traversal gives $W_{net} < 0$: the surroundings do net work on the gas over the cycle. The sign is therefore, **negative**.

(b) Over a complete cycle the gas returns to its starting state, so $\Delta U_{cycle} = 0$ and the first law reduces to $Q_{net} = W_{net}$. Since $W_{net} < 0$, the net heat absorbed by the gas is also **negative**; on balance, heat flows out of the gas to the surroundings.

(c) The device consumes work and rejects net heat to the surroundings. This is the operating signature of a refrigerator (or, equivalently, a heat pump), **not** of a heat engine. (A heat engine traverses its cycle clockwise, producing positive net work and absorbing positive net heat; the counter-clockwise device does the opposite of both.)

12. A heat pump does not generate heat; it **moves** heat from one place to another, with the work input paying for the transport rather than for the heat itself. Applying the first law to the working substance over a complete cycle, $\Delta U_{cycle} = 0 = Q_{cold} - Q_{hot} + W_{in}$, where Q_{cold} is the heat absorbed from the cool exterior, Q_{hot} is the heat rejected to the warm interior, and W_{in} is the electrical work supplied. Rearranging, $Q_{hot} = Q_{cold} + W_{in}$. The heat delivered to the house is therefore the sum of the heat extracted from outside and the electrical work consumed; both contribute, and the total exceeds the work input by exactly the amount of heat extracted from the exterior. The first law is satisfied because energy is conserved: the outside cools slightly to compensate for the energy delivered to the inside. The Coefficient of Performance for heating, $COP_{heating} = Q_{hot}/W_{in}$, can therefore be considerably greater than unity (typically 3 to 5 for a domestic unit), without any violation of energy conservation.

13. The two statements are not unrelated; they are two faces of a single underlying principle, namely that there is a preferred direction for natural processes that energy conservation alone does not capture. Heat flows spontaneously from hot to cold (Clausius), and a cyclic engine cannot extract all of the heat absorbed as useful work without rejecting some to a cooler reservoir (Kelvin-Planck). Both are observed truths about the directionality of physical processes, and the two are logically equivalent: a hypothetical device that violated one could be combined with a normal device to violate the other.

For example, suppose a perfect engine existed that converted all heat absorbed from a single hot reservoir into useful work, in violation of Kelvin-Planck. Couple its work output to drive a Carnot refrigerator, which uses that work to pump heat from a cold reservoir to the same hot reservoir. The combination would extract heat from the cold reservoir, dump it into the hot reservoir, and produce no other effect, in violation of Clausius.

The construction works in reverse too, starting from a hypothetical Clausius violator. Both statements therefore stand or fall together: they express a single physical principle in two convenient phrasings, not two separate rules.

EXERCISE 5B

14. The technician confuses energy conservation, which is the first law itself, with the much stronger claim that heat input equals work input. The first law states $\Delta U = Q - W$ (with W done by the gas), or equivalently $Q = \Delta U + W$; the heat exchanged with the surroundings goes partly into raising the gas's internal energy and partly into work. During compression the surroundings do positive work on the gas (W done by gas is negative); the gas's internal energy rises and the heat exchanged is in general not equal in magnitude to the work, except in the special case of an isothermal process where the temperature happens to be held fixed.

The everyday observation that demonstrates the flaw: any air compressor at a factory or petrol station feels noticeably hot during operation, and a stream of compressed air leaving a fitting is warm to the touch. The warmth is the rise in internal energy that the work has produced; the heat exchanged with the surroundings is less than the work input by exactly this amount. If heat input and work input were always equal, the gas would never get warm during a compression, contradicting daily experience.

15. The two interpretations have opposite signs in the first-law equation, and lead to opposite conclusions about what is happening to the internal energy of the gas. Reading 'the gas doing 5000J of work' as $W_{by} = +5000\text{J}$ (work done by the gas on the surroundings) means the gas is expanding, and from $\Delta U = Q - W$ the internal energy is falling unless heat is supplied to compensate. Reading the same phrase as $W_{on} = +5000\text{J}$ (work done by the surroundings on the gas) means the gas is being compressed, and the internal energy is rising unless heat is removed. Whether the gas is heating up or cooling down therefore depends on which interpretation the reader assumes.

A careful technician records the direction explicitly: 'the gas expanded against the piston, doing 5000J of work on the piston' if the gas is the work source, or 'the piston compressed the gas, with the gas absorbing 5000J of work from the piston' if the surroundings are. Once the direction is specified, the sign of W in the chosen convention is fixed and the first-law bookkeeping follows without further ambiguity. The ambiguity is a labelling defect, not a physics one.

16. The cylinder's steel walls are rigid: its internal volume is fixed. Heating from the surrounding fire is therefore an isochoric process, and the first law at constant volume reduces to $Q = \Delta U = nC_v\Delta T$. All heat absorbed from the fire is converted into a rise in internal energy of the gas, which means a rise in temperature, which (via the equation of state $PV = nRT$ at constant V) means a proportional rise in pressure. The pressure can rise rapidly as the temperature climbs, eventually exceeding the cylinder's rated working pressure and rupturing the steel walls with explosive force. The high-pressure rating is set for normal use, not for fire conditions.

The safety feature manufacturers add is a pressure-relief valve, typically a frangible disc or a spring-loaded valve set to open well below the rupture pressure of the cylinder. As the pressure rises in the fire, the valve opens before the cylinder fails, allowing the gas to escape and ignite as a controlled flame rather than as a destructive explosion. The valve converts a potential bomb into a controlled torch, which trained firefighters can manage.

17. In an open saucepan the surface of the milk is exposed to atmospheric pressure, which is constant on the timescale of the heating. The milk therefore undergoes an isobaric thermal expansion: as its temperature rises, its volume increases at constant pressure, governed by the volumetric expansion coefficient. With the cross-sectional area of the pan fixed, the increase in volume shows up as an increase in depth, and the surface visibly rises.

Once the milk reaches its boiling point, additional heat input goes into the latent heat of vaporisation, with the temperature pinned at the boiling point rather than rising further. With the temperature no longer increasing, the isobaric thermal expansion of the bulk liquid stops, and the surface stops rising for that reason. (Boiling itself produces some apparent rise as bubbles lift the surface, but that is a different effect from thermal expansion of the liquid.)

18. The dominant effect is the adiabatic compression of the air trapped inside the pump barrel during each downstroke. The downstroke takes a small fraction of a second, too short for the air to lose appreciable heat to the metal walls, so the work done by the cyclist on the air is converted almost entirely into a rise in internal energy and hence in temperature. Successive strokes deposit heat into the metal barrel by conduction; over many strokes the barrel accumulates the deposited heat and becomes noticeably warm to the touch.

Friction in the pump seal is at most a secondary contributor for two reasons. First, the seal area is small (a thin annular contact between the piston and the cylinder wall), so the rate at which friction can generate heat is limited. Second, the relative speed at the seal is the speed of the piston, comparable to a metre per second, which is small compared with the speeds of the gas molecules; the energy dissipated per unit time at the seal is therefore much less than the energy delivered as compression work into the bulk gas.

Direct observation closes the case: the entire pump barrel becomes warm, including parts far from the seal. The friction-only explanation predicts heating localised at the seal alone, which is not what is observed. Compressed-air heating, distributed throughout the barrel and conducted outward through the cylinder wall, matches the observation; friction does not.

19. Mr. Akilikubwa would say: the four cornerstone processes are canonical idealisations, chosen for their pedagogical clarity, not for any requirement that real engines must use them equally. Real engines are built around mechanical and thermal constraints, not around a balanced curriculum, and the cycle that emerges reflects what the engineer can build.

In a four-stroke internal combustion engine, the compression and power strokes are fast: a few milliseconds each at typical operating speeds. The gas does not have time to lose heat to the cylinder walls during a single stroke, so each is well approximated by an adiabat. The intake and exhaust strokes occur at near-atmospheric pressure (set by the manifolds and the open valve to the atmosphere or exhaust system), so each is well approximated by an isobar.

Isothermal operation would require slow strokes with active heat exchange between the gas and a constant-temperature reservoir, which is impractical for an engine running at thousands of cycles per minute. An isochoric leg would require holding the piston motionless during heating; this happens briefly only at top dead centre during combustion (and is sometimes drawn as a near-vertical rise on the indicator diagram), but contributes essentially no useful work. The engine cycle is therefore structured by what the engineer can build with cams, crankshafts, and valves, not by a desire for symmetry of physics teaching.

20. The proposed device would extract heat from a single thermal reservoir (the surrounding warm atmosphere) and convert it entirely into useful mechanical work (turning the fan), with no waste heat rejected to a cooler reservoir. This is precisely what the Kelvin-Planck statement of the second law forbids: no device operating in a cycle can absorb heat from a single reservoir and convert it entirely into work. Some heat must always be rejected to a cooler reservoir, and the work output is bounded above by the Carnot efficiency between the two reservoir temperatures.

The technician's design has only one thermal reservoir available (the warm air, all at the same temperature) and so cannot operate, regardless of how the internal mechanism is constructed. The restriction is not a defect in any particular design but a fundamental limit on how heat can be

converted to work; no engineering cleverness can circumvent it. The device is impossible in principle, in the same way that perpetual motion machines of the first kind (violating energy conservation) are impossible: both belong to a category of ambitious but doomed proposals.

21. Carnot's theorem fixes an upper bound on the efficiency of any heat engine operating between two reservoirs at temperatures T_{cold} and T_{hot}:

$$\eta_{Carnot} = 1 - \frac{T_{cold}}{T_{hot}} = 1 - \frac{300K}{500K} = 0.40 \text{ or } 40\%$$

The official's proposal claims 80%, double the Carnot ceiling between these two reservoirs. By Carnot's theorem, no heat engine, however cleverly designed, can exceed this ceiling for the given reservoir temperatures, so the proposal as stated is infeasible.

To make 80% theoretically achievable, the temperature ratio T_{cold}/T_{hot} would have to fall to 0.20 or below. With $T_{cold} \approx 300K$ (set by the temperature of available cooling water in a tropical climate), this requires $T_{hot} \geq 1500K$, far above the operating temperature of typical coal boilers and beyond what conventional boiler materials can withstand.

22. As the cylinder enters the warmer kitchen, the gas inside warms from 25°C (298K) to 40°C (313K). The cylinder is sealed and rigid, so the gas undergoes an isochoric heating: at constant volume, $P/T = \text{const}$, and the new pressure is:

$$P_{kitchen} = P_{outdoor} \times \frac{313K}{298K} \approx 1.050 \times P_{outdoor}$$

So the internal pressure has risen by about 5%. The valve, calibrated assuming the cylinder is at 25°C, opens when the gas pressure exceeds the rated limit at that temperature. At 40°C, the same gas exerts a 5% higher pressure for the same amount, so the valve will open earlier than originally intended.

The calibration is therefore **conservative** in the new environment: it triggers a premature release, before the actual unsafe pressure level (which corresponds to a still-higher temperature, where the cylinder structurally fails) is reached. To use the cylinder safely indoors, the user can either accept earlier-than-needed valve activation, recalibrate the valve for kitchen temperature, or allow the cylinder to acclimatise to room temperature before any pressure-critical use. A risky calibration would be one where the valve fails to trigger early enough; conservative calibration errs in the opposite direction and is therefore safe.

23. The accountant's confusion arises from looking only at the heat balance and missing the work input that drives the compressor. The refrigeration cycle exchanges energy with three places: it absorbs heat Q_{cold} from the server room (the cold reservoir), it accepts work W from the electrical compressor, and it rejects heat Q_{hot} to the surroundings (the warm reservoir, typically the kitchen air or outside).

Applying the first law to the working substance over a complete cycle, $\Delta U_{cycle} = 0$, so the energy balance reduces to:

$$Q_{cold} + W = Q_{hot}$$

The heat dumped therefore exceeds the heat absorbed by exactly the amount of electrical work supplied to the compressor. Total energy in (heat absorbed plus electrical work) equals total energy out (heat rejected). Once the electrical input is included in the bookkeeping, energy conservation is intact; the engineer's claim is correct, and is in fact the operating principle of every refrigeration system. The accountant should be reassured that no perpetual motion is at play; the electricity bill is paying for the difference.

24. (a) The change in internal energy is **zero**. For an ideal gas the internal energy depends only on temperature, $U = U(T)$; with the temperature held constant at 300K, $\Delta U = 0$ regardless of how the volume or pressure change during the compression.

(b) From the first law, $\Delta U = Q - W = 0$, hence $Q = W$ where W is the work done by the gas. During compression W is negative because $dV < 0$, so Q is also negative; heat must leave the gas. Physically, the surroundings do positive work on the gas as they compress it, which would otherwise raise the internal energy and hence the temperature. To keep the temperature at 300K, exactly the same amount of energy must leave the gas as heat to the surroundings. The work input is conducted away as heat, leaving the internal energy unchanged.

(c) The work done by the gas during an isothermal compression is:

$$W = nRT\ln\frac{V_f}{V_i} = 2.0mol \times 8.314Jmol^{-1}K^{-1} \times 300K \times \ln\frac{1.0L}{4.0L}$$

$$W = 4988J \times (-1.386) = -6915J$$

So the magnitude of the heat removed from the gas is $|Q| = |W| \approx 6.9kJ$, in agreement with the technician's observation. The work added to the gas (+6915J by the surroundings) is exactly compensated by the heat removed (−6915J), leaving $\Delta U = 0$.

25. Water boils when its saturation vapour pressure equals the pressure of the gas above its surface. At standard atmospheric pressure (1 atm), the saturation pressure of water reaches 1 atm at 100°C, so 100°C is the open-pan boiling point. In a sealed pressure cooker, steam released by initial boiling is trapped above the water, and as more steam accumulates the gas pressure inside the cooker rises above atmospheric, typically to about 2 atm. The boiling point of water at 2 atm is approximately 120°C, so the water in the cooker now boils at this higher temperature. Food in contact with 120°C steam cooks much more rapidly than food in contact with 100°C steam, because the rate of biochemical softening reactions roughly doubles for every 10K rise in temperature.

The safety valve role: the valve is set to open at a calibrated pressure threshold, typically about 2 atm. When the internal pressure reaches this threshold, the valve opens and steam escapes, removing energy from the cooker faster than the heater can supply it; the pressure (and therefore the boiling point) is held near the threshold value rather than rising indefinitely. Without the valve, continued heating would push the pressure and temperature above the safe limit, eventually rupturing the cooker.

EXERCISE 5C

26. (a) Path 1, after isothermal: $(P_0/2, 2V_0, T_0)$. After isochoric to $P = P_0$: $T = T_0 \times P_0/(P_0/2) = 2T_0$, so end state $(P_0, 2V_0, 2T_0)$.

Path 2, isobaric to $V = 2V_0$: $T = 2T_0$, end state $(P_0, 2V_0, 2T_0)$.

Both paths terminate at the same state B.

(b) For a monatomic gas, $\Delta U = \frac{3}{2}nR\Delta T = \frac{3}{2}nR(2T_0 - T_0) = \frac{3}{2}P_0V_0$, using $nRT_0 = P_0V_0$. Same value along both paths since U is a state function.

(c) Path 1: $W_1 = nRT_0\ln 2 + 0 = P_0V_0\ln 2$.

Path 2: $W_2 = P_0(2V_0 - V_0) = P_0V_0$.

$Q_1 = \Delta U + W_1 = \frac{3}{2}P_0V_0 + P_0V_0\ln 2$.

$Q_2 = \Delta U + W_2 = \frac{3}{2}P_0V_0 + P_0V_0 = \frac{5}{2}P_0V_0$.

$$Q_2 - Q_1 = P_0V_0(1 - \ln 2) \approx 0.307P_0V_0$$

27. $f = 7$. Molar mass $16gmol^{-1}$ identifies the gas as methane, CH_4. Methane is a non-linear polyatomic molecule with 3 translational and 3 rotational degrees of freedom (6), with the additional fractional degree reflecting partial activation of vibrational modes at room temperature, consistent with $f \approx 7$ and $\gamma = 1.30$.

28.

$$P_f = P_0\left(\frac{V_0}{V_f}\right)^{\gamma} = P_0 \times 4^{1.4} = 6.96P_0$$

$$T_f = T_0\left(\frac{V_0}{V_f}\right)^{\gamma-1} = T_0 \times 4^{0.4} = 1.74T_0$$

$$W_{on} = nC_v(T_f - T_0) = \frac{5}{2}nR(0.74T_0) = \frac{5}{2}(0.74)P_0V_0 = 1.85P_0V_0$$

29.

$$\frac{dQ}{dt} = \frac{kA\Delta T}{L} = \frac{230\text{Wm}^{-1}\text{K}^{-1} \times 0.005\text{m}^2 \times 100\text{K}}{0.005\text{m}} = 23000\text{W}$$

$$U = mc\Delta T = 5.0 \times 10^{-4}\text{kg} \times 1000\text{Jkg}^{-1}\text{K}^{-1} \times 100\text{K} = 50\text{J}$$

$$t_{cond} = \frac{U}{dQ/dt} = \frac{50\text{J}}{23000\text{W}} = 2.2\text{ms}$$

The compression stroke (1ms) is shorter than the conduction timescale (2.2ms) by a factor of about two. The adiabatic approximation is appropriate to first order, but with a non-negligible correction; most of the heat does not have time to leak during the stroke, but a small fraction does.

30. $W_{net} = 900\text{J}$, $Q_{in} = 3450\text{J}$, $\eta = 26.1\%$, the actual efficiency 26% is well below the Carnot ceiling 93.8%, as expected for a non-Carnot cycle.

31.

$$T_f = T_i\left(\frac{V_i}{V_f}\right)^{\gamma-1} = 300\text{K} \times 16^{0.4} = 909\text{K} = 636°\text{C}$$

636°C is well above the autoignition temperature 480°C, so diesel fuel injected into the compressed air ignites spontaneously without an external spark. Petrol engines use lower compression ratios (typically 10-12), so the air temperature at the end of compression remains below the autoignition point of petrol-air mixture, and a spark plug is required to initiate combustion.

32. 28.8cmHg, $-45.6°\text{C}$

33.

$$\eta_1 = 1 - \frac{T_c}{T_h} = 1 - \frac{300}{500} = 0.40$$

$$Q_h = \frac{W_1}{\eta_1} = \frac{200\text{J}}{0.40} = 500\text{J}$$

$$\eta_2 = 1 - \frac{250}{500} = 0.50$$

$$W_2 = \eta_2 \times Q_h = 0.50 \times 500\text{J} = 250\text{J}$$

$$\frac{\eta_2 - \eta_1}{\eta_1} = \frac{0.10}{0.40} = 25\%$$

New work output is 250J per cycle; the percentage increase in efficiency is 25%.

34.

$$n = \frac{P_1V_1}{RT_1} = \frac{1.0 \times 10^5\text{Pa} \times 0.020\text{m}^3}{8.314\text{Jmol}^{-1}\text{K}^{-1} \times 300\text{K}} = 0.802\text{mol}$$

Stage 1: adiabatic compression to $V_2 = V_1/4 = 5.0 \times 10^{-3}\text{m}^3$.

$$T_2 = T_1 \times 4^{\gamma-1} = 300\text{K} \times 4^{0.4} = 522\text{K}$$

$$P_2 = P_1 \times 4^{\gamma} = 1.0 \times 10^5\text{Pa} \times 4^{1.4} = 6.96 \times 10^5\text{Pa}$$

$$W_{1,on} = nC_v(T_2 - T_1) = 0.802\text{mol} \times \frac{5}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times 222\text{K} \approx 3700\text{J}$$

Stage 2: isobaric cooling at P_2 from $T_2 = 522\text{K}$ to $T_3 = 300\text{K}$.

$$V_3 = V_2 \times \frac{T_3}{T_2} = 5.0 \times 10^{-3}\text{m}^3 \times \frac{300\text{K}}{522\text{K}} = 2.87 \times 10^{-3}\text{m}^3$$

$$W_{2,on} = -P_2(V_3 - V_2) = -6.96 \times 10^5\text{Pa} \times (-2.13 \times 10^{-3}\text{m}^3) \approx 1480\text{J}$$

$$W_{on,total} = W_{1,on} + W_{2,on} = 3700\text{J} + 1480\text{J} = 5180\text{J}$$

$$Q_2 = nC_p(T_3 - T_2) = 0.802\text{mol} \times \frac{7}{2} \times 8.314\text{Jmol}^{-1}\text{K}^{-1} \times (-222\text{K}) \approx -5180\text{J}$$

$$\Delta U_{total} = nC_v(T_3 - T_1) = 0$$

Total work done on the gas is approximately 5180J, the heat rejected during the cooling stage is approximately 5180J, and the net change in internal energy is zero.

35.

Before inversion (air below mercury): $P_{air} = P_{atm} + h_{Hg} = 76\text{cmHg} + 20\text{cmHg} = 96\text{cmHg}$, $L_1 = 40\text{cm}$.

After inversion (air above mercury): $P_{air} = P_{atm} - h_{Hg} = 76\text{cmHg} - 20\text{cmHg} = 56\text{cmHg}$.

Boyle's law (isothermal): $P_1L_1 = P_2L_2$.

$$L_2 = \frac{P_1L_1}{P_2} = \frac{96\text{cmHg} \times 40\text{cm}}{56\text{cmHg}} = 68.6\text{cm}$$

The trapped air column has lengthened from 40cm to approximately 68.6cm.

EXERCISE 6A

1. The condition for SHM is that the acceleration is directed towards a fixed point and proportional in magnitude to the displacement from that point, $a = -\omega^2x$. A particle in uniform circular motion has a centripetal acceleration of **constant** magnitude ω^2A directed towards the centre, so the magnitude of the acceleration does not vary with the displacement; the motion is therefore not SHM. The projection of the position vector on a

diameter, however, has displacement $x = A\cos(\omega t)$ and acceleration $\frac{d^2x}{dt^2} = -\omega^2 A\cos(\omega t) = -\omega^2 x$, which does satisfy the SHM condition. Kipute is correct.

2. The error is to confuse a free choice with a vacuous one. The experimenter is free to choose where $t = 0$ lies, but once that choice is made, ε is fixed and is no longer at the experimenter's disposal. It encodes where the oscillator is, and which way it is moving, at that chosen instant: $y(0) = A\sin\varepsilon$ and $v(0) = A\omega\cos\varepsilon$. When two oscillators of the same A and ω are observed together, their phase difference $\Delta\varepsilon = \varepsilon_1 - \varepsilon_2$ is the physical relationship between them, telling whether they are in step ($\Delta\varepsilon = 0$), out of step ($\Delta\varepsilon = \pi$), or in quadrature ($\Delta\varepsilon = \pi/2$).

3. The two formulas describe the same situation written in two languages. Starting from $y = A\sin(\omega t + \varepsilon)$ and $v = A\omega\cos(\omega t + \varepsilon)$, squaring and adding gives:

$$y^2 + \frac{v^2}{\omega^2} = A^2$$

so that $v = \omega\sqrt{A^2 - y^2}$: Kipanga's formula is in y and Kipute's in t. The speed is maximum at $y = 0$ because that is where the oscillator has been falling "downhill" in the energy diagram for the longest time without yet having to climb; it is zero at $y = \pm A$ because all the kinetic energy has there been converted to potential energy. The acceleration is unrelated to the speed; it depends on the displacement only, and so is greatest where the speed is least.

4. The error is the assumption that zero velocity implies zero acceleration. The acceleration in SHM is $a = -\omega^2 x$; it depends on the displacement, not on the velocity. At the positive extreme, $x = +A$, so $a = -\omega^2 A$ has the largest magnitude it ever has, directed back towards the centre. The body is momentarily stationary at the extreme precisely because it has been decelerating throughout the previous quarter cycle and is now turning round; the acceleration is therefore at its largest, not its smallest.

5. The error is the conflation of an average with an instantaneous equality. With $x = A\sin(\omega t)$ and $v = A\omega\cos(\omega t)$, the kinetic energy is $K = \frac{1}{2}kA^2\cos^2(\omega t)$ and the potential energy is $U = \frac{1}{2}kA^2\sin^2(\omega t)$. They are nowhere equal at every instant; on the contrary, when one is at its maximum the other is zero. They are equal only on the time-average, since the mean of $\sin^2(\omega t)$ and of $\cos^2(\omega t)$ over a complete cycle is each $\frac{1}{2}$, and so each averages to $\frac{1}{4}kA^2$, together accounting for the total energy $\frac{1}{2}kA^2$. The equipartition of average energies between the two stores reflects the symmetric way in which the oscillator passes between them.

6. The error is to confuse "plays no part" with "does not appear in the period formula". Gravity does act on the mass, and it is responsible for the new equilibrium position being lower than the natural length of the spring by $e = mg/k$. When the mass is displaced by x from this new equilibrium, the net force is $F = -k(e + x) + mg = -kx$, because the constant gravitational pull mg is exactly cancelled by the constant extra spring tension ke. Gravity has therefore been absorbed into the choice of equilibrium and disappears from the equation of motion, leaving $T = 2\pi\sqrt{\frac{m}{k}}$. A weaker g simply moves the equilibrium up; the oscillation about it is unchanged.

7. The exact restoring acceleration is $a = -g\sin\theta$, which is not proportional to θ; the motion is therefore not strictly SHM. The formula $T = 2\pi\sqrt{\frac{l}{g}}$ is obtained by replacing $\sin\theta$ with θ, an approximation which is only valid when θ is small. For larger amplitudes $\sin\theta < \theta$, the restoring acceleration is weaker than the SHM model predicts, and the period is correspondingly longer; it grows with amplitude. Kipute is therefore right. Kipanga's claim is acceptable for amplitudes up to about 10°, where $\sin\theta$ differs from θ by less than 0.5%, and the resulting error in T is well below 1%.

8. The squared period rearranges as:

$$T^2 = \frac{4\pi^2(k^2 + h^2)}{gh} = \frac{4\pi^2}{g}\left(\frac{k^2}{h} + h\right),$$

exposing two competing terms that depend on h in opposite ways. As $h \to 0$, the term $\frac{k^2}{h}$ diverges and $T \to \infty$: pivoting through the centre of mass leaves no gravitational restoring torque. As h approaches the end of the stick, the term h dominates and the period rises towards $2\pi\sqrt{h/g}$. A minimum therefore lies between.

9. The shock absorber is designed close to critical damping, where the damping coefficient γ is approximately equal to ω_0; the suspension returns to equilibrium in the shortest possible time without overshooting, and a passenger feels no oscillation after the bump. The swing door is under-damped, with $\gamma < \omega_0$; the door oscillates about its closed position with a slowly decaying amplitude before friction in the hinge brings it to rest. An over-damped system would correspond to a swing door that drags shut so slowly that it never visibly oscillates and takes a long time to settle, like a heavy fire door fitted with an over-stiff closer. Each engineering choice serves a different goal: passenger comfort calls for critical damping; gentle automatic closure tolerates light under-damping; safety doors rely on heavy over-damping.

EXERCISE 6B

10. A new shock absorber is critically damped: the damping is just sufficient to bring the suspension back to its rest height in the shortest possible time, without any overshoot. As the seal wears and the absorber loses oil, the damping weakens, and the suspension enters the under-damped regime, in which it oscillates about its rest height with a slowly decaying amplitude before coming to rest. The two or three observed bobs after each bump are these under-damped oscillations.

11. After being struck, the bell oscillates at its natural frequency in the under-damped regime. Energy is steadily lost to internal friction in the metal and to the surrounding air, so the amplitude (and hence the loudness of the note) decays steadily until the sound is no longer audible. A heavier bell of the same shape and material stores more vibrational energy at any given amplitude (more mass is moving), while the rate at which energy leaves the bell is set mainly by the surface in contact with the air and is roughly the same. More energy stored at a comparable rate of loss means a longer audible decay, so the heavier bell rings for longer.

12. The floor-machine system has a particular natural frequency at which it would oscillate if struck once and left alone; this frequency is set by the mass of the machine and the stiffness of the floor's response. The slightly unbalanced rotating drum applies a small periodic force at a forcing frequency equal to the drum's rotation rate. As the spin-up sweeps the drum's frequency through the natural frequency of the floor-machine system, the small driving force arrives in step with the system's preferred rhythm, and the response builds up into resonance: the floor shakes violently. At every other speed the driving frequency is well off the natural frequency, the drive and the response fall out of step, and the shaking is small.

13. Let the expansion joints be a regular distance s apart along the road. A vehicle travelling at speed v crosses joints at the rate $f = v/s$ jolts per second, and this is the forcing frequency applied to the suspension. At $v \approx 80$km/h this rate coincides with the natural frequency of the suspension;

the small periodic jolts arrive in step with the suspension's preferred rhythm and the response builds up into resonance, giving the worst shaking. At 70 and 90km/h the forcing rate falls off the natural frequency, jolts and swing fall out of step, and the response is small.

14. Soldiers marching in step apply a periodic vertical force to the bridge at the frequency of their cadence. If this forcing frequency happens to coincide with the natural frequency of the bridge, the bridge is driven into resonance: each footfall arrives in step with the bridge's own swing, and the response builds up into a large amplitude that may damage the structure. Breaking step destroys the periodicity of the force; the bridge then sees only a series of independent (and on average, cancelling) impulses, never a sustained drive at its natural frequency, so the resonance cannot build up.

15. From $T = 2\pi\sqrt{\frac{l}{g}}$, the period grows as g decreases. At 3000m altitude g is smaller than at sea level (the body is farther from the Earth's centre), so T is larger and each "tick" takes longer; the clock therefore loses time. To restore correct timekeeping, the watchmaker must shorten the pendulum (raise the bob along the rod), since T depends on $\sqrt{l}$, and a smaller l exactly compensates a smaller g.

16. From $T = 2\pi\sqrt{\frac{l}{g}}$, a clock that runs fast has too short a period and therefore too short an effective length. The bob must be moved **down** along the rod, increasing l (and hence T) back to its correct value. g is fixed at the laboratory's location and is not under the technician's control.

17. The pan, together with the tomatoes on it, behaves as a mass m suspended from a spring of stiffness k (the scale's internal spring). When jostled, the pan is displaced from its equilibrium reading; the spring pulls it back, the pan overshoots, the spring pushes it forward, and the system executes SHM about the equilibrium reading with period,

$$T = 2\pi\sqrt{\frac{m}{k}}$$

The pointer's brief oscillation is exactly this SHM, and its period depends only on the combined pan-plus-tomato mass and on the scale's spring constant.

18. From $T = 2\pi\sqrt{\frac{m}{k}}$, a softer mattress (smaller k) has a longer natural period: each oscillation takes a long time, so the guest perceives the slow rise and fall plainly. A firmer mattress (larger k) has a much shorter period: the oscillation is so rapid that the guest barely notices it before the motion fades. The firmer mattress therefore has the higher spring constant k.

19. The mass per unit length μ of the line is the same in both cases (it depends on the line's material and thickness). When the line is tight, T is large and $v = \sqrt{T/\mu}$ is large, so the pulse crosses the line in a fraction of a second and is barely visible. When the line sags, T is small and v is correspondingly small, so the pulse travels slowly enough to be seen with the naked eye.

20. For small displacements from rest, a slender elastic branch supplies a restoring force on the mango that is approximately proportional to the displacement, $F \approx -kx$, where k depends on the stiffness of the branch. By Newton's second law $a = F/m = -(k/m)x = -\omega^2 x$ with $\omega^2 = k/m$, which is the SHM condition. The motion is therefore approximately SHM at small amplitude. The approximation fails when the displacement becomes large enough that the branch no longer behaves elastically: the wood may bend irreversibly, or the restoring force may cease to be linear in x, and the motion departs from SHM.

EXERCISE 6C

21. The static extension is $e = 0.65\text{m} - 0.50\text{m} = 0.15\text{m}$, giving $k = mg/e = (2.5\text{kg})(9.8\text{m/s}^2)/(0.15\text{m}) = 163\text{N/m}$ and $\omega = \sqrt{k/m} = \sqrt{163/2.5}\text{rad/s} = 8.08\text{rad/s}$. Take downward positive. The block is released from rest at $x(0) = +0.05\text{m}$, so the amplitude is $A = 0.05\text{m}$, and the general SHM equation is:

$$x(t) = A\sin(\omega t + \varphi)$$

Imposing the initial conditions $x(0) = +A$ and $v(0) = 0$:

$$A\sin(\varphi) = A \quad \Rightarrow \quad \sin(\varphi) = 1 \quad \Rightarrow \quad \varphi = \pi/2$$

so $x(t) = 0.05\sin(\omega t + \pi/2)\text{m}$. At $t = 1.0\text{s}$, $\omega t = 8.08\text{rad}$, and

$$x(1.0) = 0.05\sin(8.08 + \pi/2)\text{m} = 0.05\sin(9.65\text{rad})\text{m} = 0.05 \times (-0.227)\text{m} = -0.0114\text{m}$$

That is, about **1.14cm above the equilibrium position.**

Alternative solution. Recognising that the block starts at a turning point, where sine reaches its peak with phase $\pi/2$, one may write the motion directly in cosine form,

$$x(t) = A\cos(\omega t) = 0.05\cos(\omega t)\text{m}$$

which automatically satisfies $x(0) = +A$ and $v(0) = 0$. At $t = 1.0\text{s}$:

$$x(1.0) = 0.05\cos(8.08\text{rad})\text{m} = 0.05 \times (-0.227)\text{m} = -0.0114\text{m}$$

The same answer; cos and sin describe the same physical motion, with the phase constant absorbing the difference.

$\mathbf{K_{max}} = \frac{1}{2}kA^2 = \frac{1}{2}(163\text{N/m})(0.05\text{m})^2 = \mathbf{0.204J}$

22. 9.79m/s^2

23. With $\omega = 2\pi f = 4\pi\text{rad/s}$, the energy condition $K = U$ requires $U = E/2$, i.e. $\frac{1}{2}kx^2 = \frac{1}{2} \times \frac{1}{2}kA^2$, so

$$x = \pm A/\sqrt{2} = \pm 0.04\text{m}/\sqrt{2} = \pm 0.0283\text{m}$$

Taking $x = A\sin(\omega t)$ with the particle moving in the positive direction at $x = 0$, the time to reach $x = A/2$ satisfies $\sin(\omega t) = 1/2$, i.e. $\omega t = \pi/6$:

$$t = (\pi/6)/\omega = (\pi/6)/(4\pi\text{rad/s}) = \frac{1}{24}\text{s} = 0.0417\text{s}$$

24. 0.568s

25. It loses about 7.8s per day. The clock runs slow.

26. Treating the wire as a spring of stiffness $k = AY/L$:

$$k = \frac{(0.50 \times 10^{-6}\text{m}^2)(2.0 \times 10^{11}\text{N/m}^2)}{2.0\text{m}} = 5.0 \times 10^4\text{N/m}$$

Period of vertical oscillations:

$$T = 2\pi\sqrt{\frac{m}{k}} = 2\pi\sqrt{\frac{4.0\text{kg}}{5.0 \times 10^4\text{N/m}}} = 0.0562\text{s}$$

27. 4.5N

28. From $x = 0.05\sin(4\pi t + \pi/3)$, the angular frequency is $\omega = 4\pi$rad/s and the amplitude is A = 0.05m. The first positive time at which x = +0.025m satisfies $\sin(4\pi t + \pi/3) = 0.5$. The first positive root is $4\pi t + \pi/3 = 5\pi/6$, so $4\pi t = \pi/2$ and

$$t = 1/8\text{s} = 0.125\text{s}$$

The velocity at this instant is

$$v = A\omega\cos(5\pi/6) = (0.05\text{m})(4\pi\text{rad/s})(-\frac{\sqrt{3}}{2}) = -0.544\text{m/s}$$

The total mechanical energy is

$$E = \frac{1}{2}m\omega^2A^2 = \frac{1}{2}(0.20\text{kg})(4\pi\text{rad/s})^2(0.05\text{m})^2 = 0.0395\text{J}$$

29. 22.1s

30. KE = 0.0800J, x = ±0.0577m

31. The amplitude is half the peak-to-peak range, A = 0.60m, and the angular frequency is $\omega = 2\pi/T = \pi/2$rad/s = 1.57rad/s. The maximum vertical acceleration is

$$a_{max} = \omega^2A = (1.57\text{rad/s})^2(0.60\text{m}) = 1.48\text{m/s}^2$$

Taking $x = -A\cos(\omega t)$ so the boat is at its lowest point at t = 0, the position 0.30m above the mean is reached when $\cos(\omega t) = -1/2$, i.e. $\omega t = 2\pi/3$:

$$t = (2\pi/3)/(\pi/2)\text{s} = 4/3\text{s} = 1.33\text{s}$$

32. Mass per unit length: $\mu = M/L = (1.2 \times 10^{-3}\text{kg})/(0.60\text{m}) = 2.0 \times 10^{-3}\text{kg/m}$.

Wave speed: $v = \sqrt{T/\mu} = \sqrt{80/0.002}\text{m/s} = 200\text{m/s}$. The round trip is twice the wire length:

$$t_{rt} = 2L/v = (2 \times 0.60\text{m})/(200\text{m/s}) = 6.0 \times 10^{-3}\text{s} = 6.0\text{ms}$$

With μ doubled at fixed tension, v falls by $\sqrt{2}$ to $v_{new} = v/\sqrt{2} = 141$m/s, and the new round-trip time is

$$t_{rt,new} = 2L/v_{new} = (1.20\text{m})/(141\text{m/s}) = 8.49 \times 10^{-3}\text{s} = 8.5\text{ms}$$

EXERCISE 7A

1. A progressive wave is a propagating disturbance whose speed is determined entirely by the medium's elastic and inertial properties: the elastic stiffness of the medium provides the restoring force that passes the disturbance from one particle to the next, and the inertia of those particles fixes how quickly the restoring force can accelerate them. The amplitude of the wave is set independently by the source. Increasing the amplitude makes each particle oscillate over a wider range, but it does not change the medium's stiffness or its inertia, so the rate at which the disturbance is handed on from particle to particle, and therefore the wave speed, is unaffected.

2. The two formulas are equivalent because the angular frequency and the linear frequency are connected by the definition $\omega = 2\pi f$. Substituting into Kipanga's expression gives:

$$v_{p,max} = A\omega = A(2\pi f) = 2\pi fA,$$

which is identical to Kipute's expression. Both quantities describe the maximum speed at which a single particle of the medium oscillates as the wave passes, and both are written purely in terms of the source amplitude and the rate of oscillation.

3. The source on one side of the boundary is producing oscillations at a fixed rate, equal to its frequency. Continuity of the medium at the boundary requires the same rate of oscillations to emerge on the far side: oscillations cannot appear from nowhere, and they cannot disappear without trace. The number of crests crossing the boundary per second is the same on both sides. The frequency is therefore identical on both sides, even though the wave speed and the wavelength generally differ between the two media.

4. The acoustic power P of the source is constant, and the wavefronts spread outward over expanding spheres of surface area $4\pi r^2$. The intensity is therefore:

$$I = \frac{P}{4\pi r^2},$$

so $I \propto 1/r^2$. The intensity of a progressive wave is also proportional to the square of the displacement amplitude:

$$I \propto A^2$$

Combining the two:

$$A^2 \propto \frac{1}{r^2}, \qquad \text{hence} \qquad A \propto \frac{1}{r}$$

The displacement amplitude therefore falls as $1/r$.

5. Energy is not destroyed at a silent point; it is redistributed away from it to the regions of constructive interference. Across any small region of the field around the two sources, silent zones (destructive interference) are accompanied by louder zones (constructive interference) a fraction of a wavelength away. Integrating the power flowing through any closed surface enclosing the two sources gives a total equal to the power emitted by both sources together, as required by conservation of energy. The "missing" energy at a silent point is exactly accounted for by the "extra" energy at a nearby loud point.

6. The formula $A_R = 2A\cos(\Delta\varphi/2)$ comes from adding $y_1 = A\sin(\omega t)$ and $y_2 = A\sin(\omega t + \Delta\varphi)$ by the sum-to-product identity:

$$\sin\alpha + \sin\beta = 2\sin\frac{\alpha+\beta}{2}\cos\frac{\alpha-\beta}{2}$$

The pre-factor of 2A arises only because the two waves share the same amplitude A; the identity collapses to a single sine of amplitude $2A\cos(\Delta\varphi/2)$. When the amplitudes are unequal, the same identity does not produce a single-sine form. The correct general expression is:

$$A_R = \sqrt{A_1^2 + A_2^2 + 2A_1A_2\cos(\Delta\varphi)},$$

which reduces to $2A\cos(\Delta\varphi/2)$ exactly when $A_1 = A_2 = A$.

7. A travelling wave has its space and time variables in a single combined argument $\omega t \pm kx$, so that the entire pattern slides rigidly through the medium.

$y_1 = A\sin(\omega t - kx)$: the argument is $\omega t - kx$, so the wave is travelling in the $+x$ direction.

$y_3 = A\sin(\omega t + kx)$: the argument is $\omega t + kx$, so the wave is travelling in the $-x$ direction.

A stationary wave has the space and time variables factorised into separate trigonometric functions, so that the spatial pattern is fixed in space while the whole pattern oscillates uniformly in time. $y_2 = 2A\sin(kx)\cos(\omega t)$ has $\sin(kx)$ multiplied by $\cos(\omega t)$, so it is a stationary wave.

8. Both students are partly right. A stationary wave does not transport any net energy along the string, as Kipanga says. But the medium itself is in motion: every particle of the string between two consecutive nodes oscillates vertically with simple harmonic motion, with an amplitude that varies sinusoidally with position (zero at the nodes, maximum at the antinodes). Each oscillating particle carries kinetic energy, with a maximum at the antinodes and zero at the nodes, exactly as Kipute observes. The total kinetic energy of the string oscillates twice each period, exchanged with the elastic potential energy of the strained string. The two energies sum to a constant total, stored in the standing wave but not transported.

9. The fundamental frequency of a stretched string fixed at both ends is:

$$f_1 = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$$

At constant length and tension, $f_1 \propto 1/\sqrt{\mu}$, so f_1 falls as μ rises. Physically, a string with greater linear mass density has greater inertia for the same tension-driven restoring force; the disturbance is passed more slowly from one element of the string to the next, the transverse wave speed $v = \sqrt{T/\mu}$ is reduced, and the time taken for one half-cycle round-trip ($L = \lambda/2$ in the fundamental) is therefore lengthened, giving a lower fundamental frequency.

10. The claim is **incorrect**. For a stretched string vibrating in its third harmonic, the angular frequency is $\omega_3 = 3\omega_1$, so at any antinode the maximum particle speed is:

$$v_{p,max} = A_{ant}\omega_3 = 3A_{ant}\omega_1,$$

three times the corresponding maximum at an antinode of the fundamental of the same antinode amplitude. But this applies only to particles at the antinodes. At the nodes of the third harmonic, located at $x = 0, L/3, 2L/3, L$, the particles do not move at all. Particles between a node and an antinode have amplitudes (and therefore peak speeds) smaller than the antinode value. The blanket statement that "every particle" moves three times as fast is therefore wrong.

11. The speed of a longitudinal wave through a uniform metal rod is:

$$v = \sqrt{\frac{Y}{\rho}},$$

where Y is Young's modulus and ρ is the density of the rod's material. The cross-sectional area enters the rod's mass per unit length ($\mu = \rho A$) and also enters the elastic restoring force per unit strain (proportional to YA). The two factors of A cancel exactly when the wave equation is written for a thin rod, leaving the speed dependent only on the ratio Y/ρ. A thicker rod has more inertia per metre but is also proportionally stiffer, and the two effects exactly compensate.

EXERCISE 7B

12. A wave is a propagating disturbance, not a flow of matter. The water beneath the leaf executes a small simple-harmonic up-and-down motion as each ripple passes; the molecules return to their original positions once the ripple has gone, having merely transmitted the disturbance to their neighbours. The leaf, riding on the water, shares only this local up-and-down motion. What moves outward is the energy of the disturbance and its geometric pattern; what does not move is the bulk of the water, and therefore the leaf does not travel with the ripple.

13. When the ultrasound pulse crosses the fat-muscle boundary, the source frequency is preserved at 5MHz: continuity of oscillation count requires the same number of cycles per second on both sides of the boundary, otherwise oscillations would have to appear from nowhere or disappear without trace. The wave speed, however, changes because muscle has different elastic and inertial properties from fat. Since $v = f\lambda$ with f fixed, the wavelength scales in proportion to the wave speed: the wavelength in muscle is longer than in fat by the ratio of the wave speeds. The amplitude is also reduced on transmission, because part of the wave energy is reflected at the boundary owing to the acoustic-impedance mismatch.

14. Sound travels as a longitudinal wave whose speed is set by $v = \sqrt{E/\rho}$, where E is the relevant elastic modulus of the medium. Solid soil has a much greater elastic modulus than air, while its density is only about a thousand times greater. The ratio E/ρ is therefore much larger in soil than in air, and the sound speed in soil (typically several thousand metres per second) is many times greater than in air ($340\,m/s$). A pulse generated by the cattle's hooves reaches the listener through the soil long before the same disturbance reaches him through the air, giving the earlier warning.

15. The two coherent loudspeakers, driven in phase at 680Hz, radiate sound of wavelength

$$\lambda = \frac{v}{f} = \frac{340\,m/s}{680\,Hz} = 0.50\,m.$$

As the patron walks along the line parallel to the stage, the path difference Δx from the two loudspeakers changes continuously. At positions where $\Delta x = n\lambda$ (with integer n), the two waves arrive in phase and superpose constructively, producing a loud note. At positions where $\Delta x = (n + 1/2)\lambda$, the two waves arrive in antiphase and cancel, producing a faint or silent zone. The patron walks through alternating regions of constructive and destructive interference, hearing the loudness rise and fall at regular intervals.

16. The initial beat frequency is $f_b = 6Hz$, so the string's frequency lies at $440Hz \pm 6Hz$, i.e., either 434Hz or 446Hz. Tightening a string raises its tension and therefore its frequency, by the second law of vibration $f \propto \sqrt{T}$. After tightening, the beat rate has fallen to 2Hz, which means the string's frequency is now closer to that of the fork. If the string had originally been at 446Hz, tightening would push it further above the fork and increase the beat rate (inconsistent). If it had originally been at 434Hz, tightening raises it to 438Hz and the beat rate falls to 2Hz (consistent). The original frequency of the fifth string was therefore 434Hz.

17. Pressing the string against a fret reduces its vibrating length to the section of string between the fret and the bridge. The fundamental frequency of a stretched string at constant tension and mass per unit length varies as $f_1 \propto 1/L$ (first law of vibration). To raise the pitch by a full octave, the frequency must double, so the vibrating length must be halved:

$$L_{new} = \frac{L}{2}.$$

The twelfth fret therefore lies exactly halfway along the open-string vibrating length, i.e., midway between the nut and the bridge.

18. A cable suspended between two fixed supports is mechanically equivalent to a stretched string and supports a discrete set of natural frequencies of transverse vibration, $f_n = (n/2L)\sqrt{T/\mu}$, with integer n. The wind, blowing past the cable, sheds vortices alternately on either side at a frequency proportional to the wind speed (the Strouhal frequency). At most wind speeds, this forcing frequency is detuned from every natural frequency of the cable; the cable's response is a small forced oscillation, barely visible. At one particular wind speed, the forcing frequency matches one of the cable's natural frequencies. The resonance condition is met: each shed vortex delivers energy in step with the cable's own oscillation, the amplitude builds up cycle by cycle, and the cable swings vigorously.

19. The speed of a transverse wave on a stretched string is

$$v = \sqrt{\frac{T}{\mu}},$$

where T is the tension in the string and μ is the linear mass density. Tightening the washing-line gives it a large tension and therefore a large wave speed; the pulse crosses the line in a fraction of a second. Reducing the tension to almost zero (a slack droop) leaves μ essentially unchanged but cuts v by a large factor. The same disturbance now crosses the line slowly enough that the eye can resolve the wave's progress.

20. Two tuning forks of nominally identical frequency are never exactly identical in practice. Small manufacturing differences in mass distribution, prong length, or mounting shift each fork's natural frequency slightly. When sounded together, the small frequency difference Δf produces beats: the two waves drift slowly into and out of phase at the listener's ear, giving a resultant whose amplitude rises and falls at the beat frequency

$$f_b = |f_1 - f_2|.$$

A beat rate of about 2Hz, observed by the class, implies that the two forks differ in frequency by about 2Hz despite their nominally identical 440Hz rating.

21. A long stay-cable behaves as a stretched string fixed at both ends, supporting a discrete set of natural transverse frequencies $f_n = (n/2L)\sqrt{T/\mu}$. A steady wind blowing past the cable sheds vortices alternately on either side at a frequency proportional to the wind speed. At most wind speeds this vortex-shedding frequency is detuned from every natural frequency of the cable, and the cable's response amplitude is too small to radiate audible sound. At one particular wind speed, however, the shedding frequency matches one of the cable's natural frequencies. The resonance condition is met, the amplitude of the cable's vibration builds up, and the resulting standing wave radiates an audible note. Above and below this speed the resonance is lost and the cable returns to silence.

EXERCISE 7C

22. Wavelength: 1.0m, Period: 0.040s, Wave speed: 25m/s, Maximum particle speed: 9.42m/s, Maximum particle acceleration: 1480m/s^2.

23. $I_1 = 1.27 \times 10^{-2}\text{W/m}^2$, $I_2 = 3.18 \times 10^{-3}\text{W/m}^2$, $\frac{A_1}{A_2} = 2$.

24. The resultant has amplitude 4.0cm and lags y_1 by $\pi/3$rad.

25. Wavelength: 0.50m, Path difference: 1.50m.

Since $\Delta\varphi = 6\pi = 3(2\pi)$, the path difference is exactly three full wavelengths ($\Delta x = 3\lambda$). The two waves arrive at P in phase and superpose constructively, so P lies in a region of constructive interference.

26. Increasing the tension raises the frequency of a stretched wire (second law of vibration, $f \propto \sqrt{T}$). After increasing the tension the beats vanish, so the wire is now at 512Hz. The original frequency must therefore have been below 512Hz, namely: $f_{old} = 504$Hz. $T_{new} = 92.9$N.

27. Speed: 192m/s, Tension: 184N.

28. Wave speed: 200m/s, Fundamental frequency: 83.33Hz, Highest audible harmonic: 240th

29. The wave speed on a stretched wire is $v = \sqrt{T/\mu}$, where $\mu = \rho A$ is the linear mass density of a uniform wire of cross-section A. The maximum tension that the wire can sustain is $T_{max} = \sigma_{max}A$. Substituting:

$$v_{max} = \sqrt{\frac{T_{max}}{\mu}} = \sqrt{\frac{\sigma_{max}A}{\rho A}} = \sqrt{\frac{\sigma_{max}}{\rho}}$$

The cross-sectional area cancels exactly, leaving v_{max} a function only of the material constants σ_{max} and ρ.

For copper:

$$v_{max} = \sqrt{\frac{2.2 \times 10^8\text{N/m}^2}{8900\text{kg/m}^3}} \approx 157\text{m/s}$$

Maximum tension on a wire of cross-section $A = 1.5 \times 10^{-6}\text{m}^2$:

$$T_{max} = \sigma_{max}A = (2.2 \times 10^8\text{N/m}^2)(1.5 \times 10^{-6}\text{m}^2) = 330\text{N}.$$

30. Speed: 5092m/s, Time: 9.82×10^{-3}s.

31. Amplitude: 0.015m, Wavelength: 0.50m, Wave speed: 20m/s,

Antinodes occur where $\sin(kx) = \pm 1$, i.e., $kx = (2n+1)\pi/2$ for integer $n \geq 0$. The nearest to $x = 0$ corresponds to $n = 0$:

$$x = \frac{\pi/2}{k} = \frac{\pi/2}{4\pi\text{rad/m}} = 0.125\text{m}.$$

Displacement of that antinode at $t = 0.020$s: 9.27×10^{-3}m.

EXERCISE 8A

1. A stable interference pattern requires the two illuminating sources to maintain a fixed phase relationship over a time long compared with the detector's integration time. Two physically separate tungsten lamps emit independent sequences of wavetrains, each lasting only a few nanoseconds before being replaced by the next emission. The phase relationship between the two lamps therefore shifts randomly on this nanosecond timescale, while the eye and the screen integrate over much longer times. The pattern of fringes is averaged out to a uniform illumination, because no coherence exists between the two independent sources. To make the experiment work, both beams must be derived from a single primary source (for example, by passing light from one narrow slit through two further slits).

2. For an air wedge of small angle θ between two glass plates, a dark fringe of order n in reflected light occurs where the air-gap thickness t satisfies $2t = n\lambda$ (allowing for the π phase change on reflection at the lower glass surface). The position x_n along the wedge at which this thickness is reached is:

$$x_n = \frac{t_n}{\theta} = \frac{n\lambda}{2\theta}$$

The spacing between consecutive dark fringes is therefore:

$$\Delta x = x_{n+1} - x_n = \frac{\lambda}{2\theta},$$

which depends only on the wavelength and the wedge angle and not on the position along the wedge. Consecutive dark fringes are therefore equally spaced everywhere along the wedge.

3. The fine periodic structure of the data tracks on the reflective surface of a compact disc makes the disc act as a reflection diffraction grating. For light of wavelength λ incident at angle i and diffracted at angle θ relative to the disc's normal, the principal maxima satisfy

$$d(\sin i + \sin\theta) = n\lambda,$$

with d the track spacing. For a given disc and a given source, different wavelengths satisfy this condition at different angles θ. As the disc is tilted, i changes; the diffracted-maximum angle θ for each wavelength changes in step, so the colour seen at any given viewing direction shifts continuously with the disc's orientation.

4. The student's claim is incorrect. The bright fringes occur where the path difference between the two slits equals an integer number of wavelengths, $\delta = n\lambda$, giving positions $y_{n,br} = n\lambda D/d$. The dark fringes occur where the path difference is an odd half-integer of wavelengths, $\delta = (n + 1/2)\lambda$, giving positions $y_{n,dk} = (n + 1/2)\lambda D/d$. Each dark fringe therefore lies exactly midway between the bright fringe of order n and the bright fringe of order $n + 1$, at a precise and predictable position fixed by the path-difference condition.

5. Kipute is correct. In reflected light, the two interfering beams reflect from the lower surface of the lens (no phase change, glass-to-air) and from the upper surface of the plate (a π phase change, air-to-glass). At the centre, where the geometric path difference vanishes, the two beams differ in phase by π and interfere destructively, giving a dark centre.

In transmitted light, the two interfering beams are the directly transmitted beam (with no internal reflections) and the beam that has undergone one reflection at the upper plate surface followed by one reflection at the lower lens surface. These two internal reflections occur at the same kind of interface (air into glass on each side) and contribute equal phase shifts, so the net relative phase between the directly transmitted and the doubly reflected beams is zero. At the centre, the path difference also vanishes, so the two beams arrive in phase and interfere constructively. The centre is therefore bright. The reflected and transmitted patterns are complementary, consistent with conservation of energy.

6. The student's distinction is not physically meaningful. Diffraction and interference are both manifestations of the same underlying principle: the superposition of coherent secondary wavelets according to Huygens' principle. When a wavefront passes through any aperture, each point on the wavefront acts as a secondary source, and the field at any observation point is the coherent sum of the contributions from all those secondary sources. When the aperture is small (a single slit), the secondary sources form a continuum and the resulting pattern is conventionally called a diffraction pattern. When the aperture is two separate slits, the secondary sources cluster around two locations and the resulting fine fringe structure is conventionally called an interference pattern. Both are interference patterns; the labels are historical.

7. The first student is correct. The dark fringes of a single-slit pattern occur where $a\sin\theta = n\lambda$ for non-zero integers n. The angular separation between consecutive dark fringes is λ/a, the same for every pair. The central maximum, however, is bounded on both sides by the first dark fringe at $\sin\theta = \pm\lambda/a$, so its angular full-width is $2\lambda/a$. Every secondary maximum is bounded by two consecutive dark fringes a distance λ/a apart, so it is half the width of the central maximum. The factor of two arises because the central maximum is bounded on both sides by the same diffraction order (the first), while every other maximum is bounded by two successive orders.

8. An unpolarised beam consists of light whose electric vector has no preferred direction in the plane perpendicular to the direction of propagation. The instantaneous electric vector can be decomposed into two perpendicular components: one parallel to the Polaroid's transmission axis and one perpendicular to it. Because the unpolarised light's electric vector points equally probably in any direction in the transverse plane, the two components carry, on average, equal intensity. The Polaroid transmits the component along its transmission axis and absorbs the component perpendicular to it. Exactly one half of the incident intensity is therefore transmitted. The result is independent of the orientation of the Polaroid's transmission axis in the transverse plane, because any orientation defines two perpendicular components that share the intensity equally for unpolarised light.

9. The error in Kipanga's claim is the implicit assumption that the refractive index of a transparent medium is independent of wavelength. In reality every real transparent medium is dispersive: its refractive index varies with the wavelength of the light, with shorter wavelengths typically refracting more strongly than longer ones. Brewster's angle is $\tan i_p = n$. Since n depends on λ, so does i_p. Different colours of light therefore have slightly different Brewster angles, and at any single fixed angle of incidence the polarisation is exact only for one wavelength at a time. The variation across the visible spectrum is small (typically a fraction of a degree), but it is real and is the reason that the polarisation of reflected white light is never perfect at any single angle.

10. The bright fringes in a Young's double-slit pattern occur at positions on the screen where the path difference from the two slits equals an integer number of wavelengths,

$$\delta = n\lambda \qquad \Rightarrow \qquad y_{n,br} = \frac{n\lambda D}{d}$$

The dark fringes occur at positions where the path difference equals an odd half-integer of wavelengths,

$$\delta = (n + 1/2)\lambda \qquad \Rightarrow \qquad y_{n,dk} = \frac{(n + 1/2)\lambda D}{d}$$

The position of the nth dark fringe relative to the nth bright fringe is:

$$y_{n,dk} - y_{n,br} = \frac{\lambda D}{2d} = \frac{\beta}{2},$$

where $\beta = \lambda D/d$ is the bright-fringe spacing. The dark fringe therefore lies exactly halfway between the bright fringe of order n and the bright fringe of order $n + 1$.

11. The student's error is to assume that a change in wavelength shifts the central fringe. The position of the central fringe is determined by zero path difference between the two beams, not by the wavelength. Immersing the entire apparatus in water does not break the symmetry of the two paths from the slits to the centre of the screen, because both paths travel through water of the same refractive index. The path difference at the centre therefore remains zero, and the central fringe stays where it was.

What actually happens is that the wavelength of the light inside water is reduced by a factor equal to the refractive index of water, so the fringe spacing $\beta = \lambda D/d$ shrinks by the same factor. All non-central fringes therefore move closer to the centre while the central fringe is anchored. The pattern is compressed about its centre, but it is not shifted.

12. The dark ring of order n in a Newton's-rings pattern occurs where the air-gap thickness t satisfies $2t = n\lambda$. The geometric relation between the ring radius r and the gap thickness t for a spherical lens of radius of curvature R resting on a flat plate, valid for small t, is $r^2 = 2tR$ (from the chord-of-circle approximation). Combining the two relations:

$$r^2 = 2\left(\frac{n\lambda}{2}\right)R = n\lambda R \qquad \Rightarrow \qquad r = \sqrt{n\lambda R}$$

The ring radius therefore grows as $\sqrt{n}$, not in direct proportion to n. The geometric origin is the curvature of the lens: as one moves outward from the centre, the air-gap thickness grows quadratically with radius rather than linearly, so a unit increase in ring order (corresponding to a fixed linear increase in thickness) corresponds to a smaller and smaller incremental radial step as the rings move outward. The rings therefore crowd together as n grows.

13. Both the soap film and the glass pane are thin transparent layers bounded by two reflecting surfaces, and in principle both should produce interference between the front-reflected and back-reflected beams. The difference lies in the comparison between the coherence length of the white-light illumination and the optical path difference $2nt$ across each layer.

White light from an ordinary source has a very short coherence length, typically a few micrometres. The two reflected beams interfere only if the path difference $2nt$ is smaller than this coherence length. A soap film is typically a few hundred nanometres to a few micrometres thick, so $2nt$ fits comfortably within the coherence length and the two reflections interfere to give visible colour bands.

A glass window-pane is several millimetres thick, so $2nt$ is many thousands of times the coherence length of white light. The two reflections from the front and back surfaces of the pane are then essentially independent (incoherent), and no interference fringes are produced. The pane simply transmits the white light unchanged, with two faint uncoloured reflections.

EXERCISE 8B

14. The Newton's-rings pattern in reflected light arises from interference between two beams: one reflecting from the lower surface of the lens (a glass-to-air interface, with no phase change) and one reflecting from the upper surface of the flat plate (an air-to-glass interface, with a π phase change). At the centre of the pattern, where the air-gap thickness is essentially zero, the two beams have zero geometric path difference but they differ in phase by π, so they cancel destructively and the centre is dark.

Silvering the lower surface of the lens replaces the glass-to-air interface with a perfect mirror. All incident light is now reflected back at the lens itself; no light enters the air gap to reach the plate below. The interference between the two air-gap reflections, which is the entire mechanism by which the Newton's rings are produced, is therefore destroyed. The field of view no longer shows any ring pattern, only a uniform bright reflection from the silvered surface. The previously dark centre is replaced by a uniform bright field.

15. The light emerging from a smartphone screen is plane-polarised, because the LCD's outer polariser fixes the direction of the electric vector along a single axis (typically at 45° to the screen edges). When this polarisation direction is parallel to the transmission axis of the wearer's polarised sunglasses, the light is transmitted and the screen looks bright. When the phone is rotated through 90° about the viewing axis, the LCD's polarisation direction also rotates through 90° and becomes perpendicular to the sunglasses' axis. By Malus's law,

$$I = I_0\cos^2\theta = I_0\cos^2(90^\circ) = 0,$$

so no light reaches the eye through the sunglasses, and the screen appears black. The observation is a direct consequence of the LCD emitting already-polarised light and the sunglasses acting as an analyser whose transmission depends on the relative orientation of the two axes.

16. A soap bubble is a thin film of soapy water held together by surface tension. Gravity drains the liquid downward inside the film, so that the film becomes thicker at the bottom and thinner at the top. The thickness t varies continuously and smoothly with height on the bubble's surface.

Light reflected from the front and back surfaces of the film interferes. The condition for a bright reflected fringe at wavelength λ is

$$2nt = (m + 1/2)\lambda,$$

with the extra half-wavelength accounting for the π phase change on reflection at the front surface. At each thickness t, this condition picks out a specific wavelength for bright reflection, and so a specific colour appears at that height. As the film continues to drain, t changes continuously with time at every height; the height at which any given colour satisfies the bright condition therefore moves downward, and the colour bands flow downward across the bubble's surface.

17. A thin layer of petrol on water acts as a thin film bounded by two reflecting surfaces, the air-petrol interface and the petrol-water interface. The two reflections interfere, with optical path difference $2nt\cos r$ between them, where n is the refractive index of petrol, t is the film thickness and r is the angle of refraction inside the film. The condition for a bright reflected fringe at wavelength λ is:

$$2nt\cos r = (m + 1/2)\lambda.$$

When the motorcyclist changes his viewing angle, the angle of refraction r changes; the effective path difference $2nt\cos r$ changes correspondingly, and the wavelength that satisfies the constructive condition at any given point on the puddle shifts to a different value. Different colours therefore appear at the same point on the petrol layer when viewed from different angles. The shifting colours are the angle-dependence of the same interference mechanism that produces them.

18. Light reflected from the surface of a body of water at large angles of incidence is partially polarised in the horizontal direction (the component of the electric vector parallel to the lake's surface is preferentially reflected). At the Brewster angle i_p for water, given by

$$\tan i_p = n_{water} = 1.33 \quad \Rightarrow \quad i_p \approx 53^\circ,$$

the reflected light is completely plane-polarised in the horizontal direction. Polaroid sunglasses with a vertical transmission axis block almost all of this horizontally polarised glare. The sunglasses are therefore most effective when the fisherman views the lake along a line of sight close to Brewster's angle from the vertical, i.e., his line of sight makes about 53° with the normal to the lake (about 37° above the horizontal). At other angles of view, the reflected light is only partially polarised and the suppression is less complete.

19. On a hot day, the tarmac surface heats the air immediately above it to a temperature much higher than the air higher up. Hot air is less dense than cool air, and so has a slightly smaller refractive index. The refractive index of the air therefore decreases continuously as one approaches the road surface.

A wavefront of light from the sky, approaching the road at a glancing angle, passes through layers of progressively smaller refractive index. By Snell's law, the wavefront bends further from the normal at each layer interface, and eventually undergoes total internal reflection at a layer immediately above the tarmac, curving back upward toward the driver's eye. The driver sees light from the sky as if it had been reflected from a pool of water on the road. As he approaches, the geometry of the bending changes, and the "pool" appears to recede; he can never quite reach it.

20. Light incident on the coated lens reflects partially at the air-coating interface and partially at the coating-glass interface. Both reflections occur at boundaries where the index increases (air to coating, and coating to glass), so each reflection undergoes a π phase change. The two π shifts cancel, and the relative phase of the two reflected beams is set entirely by the optical path difference $2nt$ accumulated by the wave inside the coating.

Choosing the thickness so that $2nt = \lambda/2$, that is,

$$t = \frac{\lambda}{4n},$$

makes the two reflected beams meet 180° out of phase at the air-coating boundary, and they cancel destructively. The intensity reflected at the chosen wavelength of 550nm is therefore strongly reduced. The two reflections have nearly equal amplitude because the refractive index of magnesium fluoride lies between that of air and that of glass, allowing nearly complete cancellation.

21. Sunlight scattered by air molecules into a direction making a large angle with the original beam is strongly linearly polarised, a consequence of Rayleigh scattering. The scattered electric vector must be perpendicular both to the incoming direction of polarisation and to the direction of scattering, so light scattered through approximately 90° from the original beam is most strongly polarised.

At midday, light reaching the camera from directly overhead has been scattered through approximately $90°$ from the incoming sunlight, and is therefore strongly polarised. Light from a direction close to the sun has been scattered through only small angles, and is weakly polarised. When the photographer rotates the polarising filter, she varies the angle between the filter's transmission axis and the polarisation direction of the overhead scattered light; by Malus's law $I = I_0\cos^2\theta$, the transmitted intensity ranges from full at $\theta = 0$ down to zero at $\theta = 90°$. The overhead sky therefore darkens visibly at one particular orientation and brightens again at the perpendicular orientation. The sky close to the sun, being only weakly polarised, hardly changes its brightness regardless of the filter's orientation.

22. A hologram records the interference pattern between a reference beam and the light scattered from an object, in a photosensitive emulsion. Reconstruction of the original object wavefronts requires illumination of the hologram with a beam of the same wavelength as the reference beam used during recording. The diffraction angles produced by the recorded fringe pattern follow the grating equation:

$$d\sin\theta = n\lambda$$

If the illuminating wavelength matches the recording wavelength, the diffracted wavefronts reconstruct the original object wavefronts exactly, and the image appears sharp. If the illuminating wavelength is different, the diffracted angles shift, and the reconstructed wavefronts no longer align with the original object geometry; the image becomes distorted and blurred.

The security hologram was recorded with red laser light. Under the red LED inspection lamp (wavelength close to the recording wavelength), the reconstruction is faithful and the image appears sharp and three-dimensional. Under the yellow sodium street-lamps (wavelength 589nm, very different from the red recording wavelength), the diffraction angles shift, the wavefronts misalign, and the image looks blurred and washed-out.

23. Most clear plastics are optically isotropic at rest: the refractive index inside the plastic is the same for any direction of polarisation, so plane-polarised light passing through the plastic emerges still plane-polarised, and is blocked by a crossed analyser. The view between the crossed Polaroids therefore looks uniformly dark (or shows only a faint colour pattern from residual stresses left in manufacture).

When the plastic is mechanically stressed, its optical isotropy is broken: the refractive index for light polarised along the local stress direction becomes slightly different from that for light polarised perpendicular to it. The material is then said to be stress-induced birefringent. Plane-polarised light entering the stressed region emerges elliptically polarised, with the ellipse parameters depending on the local stress. When this elliptical light passes through the crossed analyser, a fraction of the incident intensity is transmitted, the fraction depending on the local stress. Under white light, different wavelengths transmit different fractions, producing a coloured pattern that maps the stress field. Squeezing the protractor increases the stress throughout the plastic, strengthens the birefringence, and brightens the pattern; releasing the squeeze removes the stress, restores isotropy, and the pattern fades.

24. A prism separates colours by refraction: the refractive index of glass depends on the wavelength of light, so different colours bend by different angles when passing through the prism. There is a single, smoothly-varying angle of deviation that depends on wavelength.

A CD has no glass and there is no refraction involved. The reflective surface of the CD is ruled with a fine periodic structure (the data tracks), and acts as a reflection diffraction grating. White light incident on the disc is reflected at angles satisfying the grating equation

$$d\sin\theta = n\lambda,$$

with d the track spacing and n an integer order. Different wavelengths satisfy this equation at different angles, separating into a spectrum at each diffraction order. The CD therefore spreads colour by wavelength-dependent diffraction off a periodic surface structure, while the prism spreads colour by wavelength-dependent refraction through a transparent medium. The two mechanisms are physically distinct, and produce slightly different orderings of the colours in the spectrum.

25. The source of the ultrasound pulse emits oscillations at a fixed rate (here 5MHz). Continuity at the fat-muscle boundary requires the same rate of oscillations on both sides of the boundary: cycles cannot be created from nothing or be destroyed without trace as the wave crosses the interface. The frequency f is therefore preserved across the boundary, in both fat and muscle.

The wave speed v, by contrast, depends on the elastic and inertial properties of the medium, which are different in fat and in muscle. Muscle has different bulk elasticity and density from fat, and the longitudinal wave speed in muscle is larger. Since $v = f\lambda$ and f is fixed, the wavelength must scale in proportion to v:

$$\lambda_{\text{muscle}} = \frac{v_{\text{muscle}}}{f} > \lambda_{\text{fat}} = \frac{v_{\text{fat}}}{f}.$$

The wavelength inside muscle is therefore larger than the wavelength inside fat, while the frequency stays the same.

EXERCISE 8C

26. $30.7°$, 400nm, $2.00\times10^{8}\text{m/s}$.

27. $2.0\times10^{-10}\text{s}$, 6.0cm.

28. 2.64mm, $0.126°$.

29. Slit separation: 0.60mm, New fringe width: 0.775mm.

30. $3.25\times10^{-6}\text{m}$ or $3.25\ \mu\text{m}$.

31. Wedge angle:

$$\theta = \frac{t}{L} = \frac{0.040\times10^{-3}\text{m}}{8.0\times10^{-2}\text{m}} = 5.0\times10^{-4}\text{rad}.$$

Spacing between consecutive dark fringes:

$$\Delta x = \frac{\lambda}{2\theta} = \frac{589\times10^{-9}\text{m}}{2(5.0\times10^{-4})} = 5.89\times10^{-4}\text{m} = 0.589\text{mm}.$$

Total number of dark fringes from the apex to the paper end (where the air-gap thickness is 0.040mm):

$$N = \frac{2t_{\text{max}}}{\lambda} = \frac{2(0.040\times10^{-3}\text{m})}{589\times10^{-9}\text{m}} \approx 135.8.$$

Rounded to the nearest integer, approximately 136 dark fringes are observed along the wedge (orders $n = 0$ at the apex up to $n = 135$ near the paper end).

32. Smallest thickness: 85.7nm, Next-smallest thickness: 257nm.

33. $r_{10} = 2.32\text{mm}$, $R' = 1.08\text{m}$.

34. 1.35.

35. 0.245mm, $0.413°$.

36. $\theta_{\text{red}} = 25.9°$, $\theta_{\text{blue}} = 16.3°$.

EXERCISE 9A

1. The statement is incorrect because, in a sound wave, the displacement of a particle and the pressure variation are a quarter of a cycle out of phase. At its displacement extreme the particle is momentarily at rest, and the air around it is then neither compressed nor rarefied, so the pressure equals the normal atmospheric value. The pressure is greatest where the displacement is zero: as the particle passes through its equilibrium position, neighbouring particles crowd in on one side while those on the other side move away, producing the maximum compression or rarefaction. Greatest pressure therefore coincides with zero displacement, not with greatest displacement.

2. Humid air contains water vapour. A water molecule has a molar mass of 18g/mol, which is less than that of the nitrogen (28g/mol) and oxygen (32g/mol) molecules it replaces, so humid air is less dense than dry air at the same temperature and pressure. The speed of sound in a gas is $v = \sqrt{(\gamma RT/M)}$, so a smaller average molar mass M gives a larger speed. Hence sound travels slightly faster through humid air.

3. All audible frequencies travel through air at the same speed. The speed of sound $v = \sqrt{(\gamma RT/M)}$ depends only on the nature of the gas and its temperature, not on the frequency of the wave; air is a non-dispersive medium for sound. The high notes and the low notes therefore cover the distance to the listener in exactly the same time, so the chords stay together and the melody arrives undistorted, however far away the listener stands.

4. A pipe can sustain a standing wave only at its discrete set of natural frequencies, the fundamental and its harmonics, because only these fit the boundary conditions of the pipe. Gentle blowing excites the fundamental. Blowing much harder supplies enough energy to set up the next allowed mode, and the air column locks onto that higher harmonic. Since no standing wave can exist at the frequencies that lie between two harmonics, the pitch cannot rise smoothly; it jumps abruptly from one allowed mode to the next.

5. A sound level of 0dB does not mean the absence of sound. The decibel scale measures the intensity of a sound relative to a fixed reference, the threshold of hearing, $I_0 = 10^{-12}$ W/m²; the level is ten times the logarithm, to base ten, of the ratio of the sound's intensity to this reference. A level of 0dB therefore means that the sound's intensity is equal to I_0, the faintest intensity a normal human ear can just detect. A negative decibel level describes a sound whose intensity is below I_0: a real sound that exists physically but is too faint for the human ear to detect.

6. Loudness is not the same quantity as intensity. Intensity is the physical power carried by the wave per unit area; loudness is the sensation the sound produces, and the sensitivity of the ear depends strongly on frequency. The ear is most sensitive to frequencies of about 2–5kHz and much less sensitive at low frequencies. The 3kHz tone lies near the peak of the ear's sensitivity, while the 50Hz tone lies where the ear responds poorly. At equal intensity, the 3kHz tone is therefore heard as the louder of the two.

7. Polarisation is the restriction of a wave's oscillation to a single transverse direction, and it is possible only for a transverse wave, whose oscillation is already perpendicular to the direction of travel and so offers a choice of plane. Light is a transverse wave and can therefore be polarised. Sound in air is a longitudinal wave: the air oscillates along the direction in which the wave travels, so there is no transverse direction for a polariser to select. A sound wave therefore cannot be polarised.

8. The electric and magnetic fields are in phase: they reach their maximum values at the same instant and fall to zero together. This is because the two fields are not independent. A changing electric field generates the magnetic field and a changing magnetic field generates the electric field, so the two are produced together and travel together as a single wave, with no continual exchange of energy between them of the kind that occurs between the kinetic and potential energy in simple harmonic motion. Their magnitudes are linked at every instant by $E = cB$.

9. Whether radiation harms the skin depends not on the total power it delivers but on its position in the electromagnetic spectrum, that is, on its frequency. The spectrum is ordered by frequency, and the higher the frequency, the more energetically the radiation interacts with the matter it meets. Ultraviolet lies at a higher frequency than visible light, high enough to disrupt the chemistry of the cells in the skin. Visible light, of lower frequency, cannot do this however intense it is. It is therefore the ultraviolet, and not the brighter visible light, that burns the skin.

10. A moving source and a moving observer do physically different things. A moving source travels through the air as it emits, chasing the wavefronts ahead of it and drawing away from those behind, so it crowds the wavefronts closer together in front and stretches them out behind: it changes the wavelength of the sound in the air itself. A moving observer leaves the wave in the air unchanged, the wavelength is unaltered, but by moving meets the wavefronts at a faster or slower rate than a stationary observer would. Because the source changes the wave while the observer changes only the rate at which the wavefronts are met, the two motions affect the observed frequency in different ways.

11. It is better understood as an expansion of the space between the galaxy and the Earth. The distant galaxies are very nearly at rest with respect to their own surroundings; they are not driving themselves through space. They are carried apart because the space lying between them steadily expands. This is supported by Hubble's law, which shows that the recession speed of a galaxy is proportional to its distance, exactly what a uniform expansion of space produces, since the more space lies between two galaxies, the more of it stretches and the faster they separate. Objects merely flying apart through a fixed space would not give this simple proportionality.

12. The Doppler shift of light responds only to the component of a galaxy's velocity along the line of sight, the radial component, because only that component changes the distance between the source and the observer and so stretches or compresses the wavelength. A galaxy moving purely across the line of sight keeps its distance from the Earth momentarily unchanged, so it produces no shift in wavelength. A measured redshift therefore reveals only the radial part of a galaxy's motion; its motion across the sky goes undetected by the shift.

EXERCISE 9B

13. Sound of very low frequency has a very long wavelength. A wave is strongly diffracted by an obstacle only when its wavelength is comparable with, or larger than, the obstacle, so the long waves of infrasound bend around hills, trees and buildings instead of being stopped by them. Low-frequency sound is also absorbed by the air far more slowly than high-frequency sound. For both reasons, infrasound travels much greater distances than ordinary speech, whose shorter waves are more easily blocked and absorbed.

14. The deep bass of the music is low-frequency sound of long wavelength, while the consonants that carry speech are high-frequency sound of short wavelength. Long-wavelength sound passes through the wall, and diffracts around its edges and any gaps, far more readily than short-wavelength sound, which the wall reflects and absorbs much more effectively. The wall therefore lets the bass through while muffling the conversation.

15. Sound and light are both waves, but their wavelengths are vastly different. The wavelength of ordinary speech is of the order of one metre, comparable with the size of the building, so the sound wave diffracts strongly around the corner and reaches Kipanga. The wavelength of light is less than one millionth of a metre, far smaller than the building, so light diffracts only negligibly and travels in practically straight lines; the building blocks it. The same obstacle therefore stops the light while the sound bends around it.

16. The speed of sound in air increases with temperature. On a cold morning the air is cooler, so sound travels more slowly and takes longer to cover the distance to the cliff and back; the echo returns later. On a hot afternoon the warmer air carries the sound faster, so the same round trip takes less time and the echo returns sooner.

17. The air column above the water in the bucket acts as a closed pipe and resonates as the bucket fills. As the water rises, the length of this air column steadily decreases. The resonant frequency of a closed air column is inversely proportional to its length, so a shorter column resonates at a higher frequency. The pitch of the sound therefore rises steadily as the bucket fills.

18. The whine of a mosquito is a faint sound of low intensity. In a quiet bedroom the background sound is fainter still, so the whine stands well above the level at which the ear can detect it. In a busy market the loud background noise raises the level a sound must reach before the ear can pick it out, and the faint whine is masked by the much louder surrounding noise. The mosquito produces the same sound in both places; what changes is whether that sound stands out against the background.

19. The barrier blocks the straight-line path of the sound from the road to the houses, casting a sound shadow behind it. But sound has a long wavelength and diffracts: at the top edge of the barrier the waves bend over and spread downwards into the shadow region behind it. Some sound therefore always reaches the houses by diffracting over the barrier, so the noise is reduced but never completely removed.

20. A mobile phone communicates by radio waves, which are electromagnetic. The walls of the lift are made of metal, which is a conductor. Radio waves cannot pass through a conductor: the electric field of the wave drives currents in the metal, and these currents reflect and absorb the wave so that very little of it enters the enclosed space. The metal lift therefore screens the phone from the signal, and the connection is lost.

21. A lightning flash is a sudden, violent surge of electric charge. Accelerating charge radiates electromagnetic waves, and the abrupt lightning discharge sends out a brief, powerful burst of radio waves spread over a wide range of frequencies. The radio receiver picks up these bursts along with the station it is tuned to, and they are heard as the crackling and popping.

22. Every object emits electromagnetic radiation because of the motion of its particles, and a warm object such as the human body emits this radiation mainly in the infrared part of the spectrum. A thermal-imaging camera is built to detect infrared rather than visible light. It forms its picture from the infrared radiation the body itself gives out, so it needs no visible light at all and works just as well in complete darkness.

23. A television remote control sends infrared radiation, which has a very short wavelength and behaves much like visible light: it travels in straight lines and is stopped by walls, so it needs an unobstructed path to the set. The radio waves used for broadcasting have a much longer wavelength; they diffract around obstacles and pass through walls, so a radio receiver works even when it is far from the transmitter and separated from it by walls. The difference follows from the great difference in wavelength between the two kinds of wave.

24. Gamma rays lie at the highest-frequency end of the electromagnetic spectrum. They are extremely penetrating, so they pass straight through sealed packaging and through the whole body of the instrument, reaching every part of it. They also interact very energetically with matter, energetically enough to destroy the cells of bacteria and other micro-organisms. A gamma source can therefore sterilise an instrument completely without the package ever being opened.

25. The eye evolved to detect the radiation that is both most plentiful and most usable at the Earth's surface. The Sun radiates most strongly in the visible band, and the Earth's atmosphere is highly transparent to visible light while absorbing much of the radiation in the neighbouring bands. The visible band is therefore the part of the spectrum that reaches the ground in the greatest amount, so an eye sensitive to that band gives the clearest view of the surroundings. Natural selection accordingly favoured eyes tuned to visible light.

26. The monitor sends a beam of ultrasound into the body. When the ultrasound meets the moving wall of the baby's heart it is reflected, and because the reflecting surface is moving, the reflected wave returns at a slightly changed frequency: this is the Doppler effect. The monitor compares the frequency of the reflected wave with the frequency it sent out. The reflected frequency is raised while the heart wall moves towards the probe and lowered while it moves away, so the rhythmic change in frequency reveals and follows the beating of the heart.

27. A redshift means a galaxy is receding from us and a blueshift means it is approaching. The expansion of space tends to carry every galaxy away from us, but each galaxy also has its own local motion through space. The Andromeda galaxy is one of our nearest neighbours, close enough that the gravitational attraction between it and the Milky Way is drawing the two together. Over so short a distance the small recession due to the expansion is outweighed by this local motion of approach, so Andromeda is on the whole moving towards us and its light is blueshifted. The general redshift appears only for galaxies far enough away that the expansion outweighs any local motion.

EXERCISE 9C

28. The speed of sound increases by about 5.2% between the cold night and the hot afternoon.

29. 0.20

30. About 16 machines are required.

31. 0.45W

32. The sound level falls by 6.0dB.

33. 5.0×10^{-10}s or 0.50ns.

34. 3840km/s, 55Mpc

35. The Hubble time is $\frac{1}{H_0}$. Expressing H_0 in basic units:

$$H_0 = \frac{70\text{km/s}}{3.09 \times 10^{19}\text{km}} = 2.27 \times 10^{-18}\text{s}^{-1}$$

Hence the Hubble time is:

$$\frac{1}{H_0} = \frac{1}{2.27 \times 10^{-18}\text{s}^{-1}} = 4.4 \times 10^{17}\text{s}$$

Dividing by the number of seconds in a year:

$$\frac{4.4 \times 10^{17}\text{s}}{3.16 \times 10^{7}\text{s}} = 1.4 \times 10^{10}\text{ years}$$

This is a rough estimate of the age of the universe, the time the present rate of expansion would have taken to produce the universe's present size.

36. 932Hz

37. 517Hz

38. (a) The angular frequency is $\omega = 2\pi f$:

$$\omega = 2\pi \times 400\text{Hz} = 2513\text{rad/s}$$

The maximum speed of the air particles is:

$$v_{max} = \omega s_0 = 2513\text{rad/s} \times 1.0 \times 10^{-5}\text{m} = 0.025\text{m/s}$$

(b) The pressure amplitude is:

$$p_0 = \rho v \omega s_0 = 1.2\text{kg/m}^3 \times 340\text{m/s} \times 2513\text{rad/s} \times 1.0 \times 10^{-5}\text{m} = 10.3\text{Pa}$$

(c) The intensity is:

$$I = \frac{{p_0}^2}{2\rho v} = \frac{(10.3\text{Pa})^2}{2 \times 1.2\text{kg/m}^3 \times 340\text{m/s}} = 0.13\text{W/m}^2$$

(d) The sound level is:

$$L = 10\log\frac{I}{I_0} = 10\log\frac{0.13\text{W/m}^2}{1.0 \times 10^{-12}\text{W/m}^2} = 111\text{dB}$$

39. 139N

40. (a) 768Hz (b) 12dB

41. 340m/s, 15°C.

42. For a radio wave with the source speed much less than c, the fractional Doppler shift is $\frac{\Delta f}{f} = \frac{v}{c}$. Hence:

$$\Delta f = \frac{fv}{c} = \frac{150 \times 10^6\text{Hz} \times 7.6 \times 10^3\text{m/s}}{3.0 \times 10^8\text{m/s}} = 3.8 \times 10^3\text{Hz}$$

The received frequency is raised by about 3.8kHz while the satellite approaches.

When the satellite passes directly overhead, its velocity is momentarily perpendicular to the line of sight. The radial component of the velocity is then zero, and since only the radial component produces a first-order Doppler shift, the shift falls to zero at that instant.

EXERCISE 10A

1. Electric field is a vector quantity, defined as the force per unit positive charge that a test charge would experience at the point. Forces are themselves vectors; the force a test charge experiences from several source charges is the vector sum of the individual forces, so the field, which is the force per unit charge, also adds as a vector. Electric potential is a scalar quantity, defined as the work done per unit positive charge in bringing a test charge from infinity to the point. Work is a scalar, the contributions to the total work from each source charge depend only on the radial distance from that source and not on direction, and scalar quantities add by ordinary arithmetic. The two quantities are combined by different rules because the underlying physical quantities they record (force and work) are themselves of different kinds, vector and scalar respectively.

2. An electric field is a region of space in which a stationary electric charge experiences a force of electrical origin. The field at a point is a vector quantity whose magnitude is the force per unit positive charge that a small positive test charge would experience at the point, and whose direction is the direction of that force. Operationally, the strength of the field at a point is measured by placing a small positive test charge q at the point, measuring the electrostatic force F on it, and computing $E = F/q$. The test charge is required to be small enough that its own field does not appreciably disturb the source charges responsible for the field being measured.

3. Consider two conducting spheres of different radii joined by a long thin wire, so that they share the same potential. The potential at the surface of each sphere is $V = kQ/R$; for a given V the charge held by a sphere is therefore proportional to its radius, $Q \propto R$. The surface area of a sphere is $4\pi R^2$, so the surface charge density is $\sigma = Q/(4\pi R^2) \propto 1/R$: the smaller the radius, the higher the surface charge density. An irregular conductor in electrostatic equilibrium behaves locally like a small portion of a sphere whose radius equals the radius of curvature at that point. Regions of high curvature (sharp points and edges) have small effective radius and therefore concentrate more charge per unit area than regions of low curvature, which behave like large-radius spheres.

4. In electrostatic equilibrium all excess charge on a conductor resides on its outer surface, and no charge is present in the bulk or in an empty cavity within it. The conducting material that bounds the cavity is at a single potential (the conductor is itself an equipotential), and the cavity is enclosed by a surface on which the potential is uniform. Apply Gauss's law to any closed surface drawn entirely within the empty cavity: the surface encloses no charge, so the total flux through it is zero. Combined with the symmetry imposed by the equipotential boundary, this forces the field at every point of the cavity to be zero. Equivalently, a non-zero field inside the cavity would drive a steady current in the surrounding conductor, in contradiction with the equilibrium condition.

5. Air, like any insulator, can withstand only a limited electric field before its molecules ionise and the air becomes momentarily conducting; this limiting field is the dielectric strength of air, approximately 3.0×10^6 V/m at ordinary conditions. Friction between the soles of the shoes and the dry carpet transfers electrons by the usual triboelectric mechanism, and the person walks away with a net charge that raises their potential to several thousand volts. As the hand approaches the metal doorknob, the gap between the fingertip and the knob shrinks, and the field across the gap rises sharply. When the field exceeds the dielectric strength of the intervening air, the air ionises, the gap becomes momentarily conducting, and the accumulated charge crosses the gap as a visible spark with an audible click. The spark is therefore the very mechanism by which the dielectric strength of air is exceeded and then restored.

6. Gauss's law states that the total electric flux through any closed surface in space is equal to the net charge enclosed by the surface, divided by the permittivity of free space: $\Phi = \oint E \cdot dA = \frac{Q_{enc}}{\varepsilon_0}$. The law is particularly useful for charge distributions of high symmetry (spherical, cylindrical, or planar) because for such distributions a Gaussian surface can be chosen on which the field magnitude is constant and either parallel or perpendicular to every element of the surface. The integral on the left then reduces to a simple product $E \times A$ and the field is found by a single line of algebra, without any need to evaluate Coulomb's-law integrals over the source distribution.

7. An equipotential surface is a surface in space all of whose points lie at the same electric potential. Three of its properties are: (i) no work is done by the electric field in moving a test charge from one point of the surface to any other, because the two points are at the same potential and the work done equals the charge times the potential difference; (ii) the electric field at every point of an equipotential surface is perpendicular to the surface, because any component of the field along the surface would do work as the charge moved, contradicting the first property; (iii) no two distinct equipotential surfaces ever intersect, because at the point of intersection the potential would simultaneously have two different values, which is impossible.

8. Argument by potential difference: the work done by the electric field in moving a charge q from a point at potential V_1 to a point at potential V_2 is $W = q(V_1 - V_2)$. On an equipotential surface every point has the same potential by definition, so $V_1 = V_2$, the potential difference is zero, and the work done is zero. Argument by directions: the electric field at every point of an equipotential surface is perpendicular to the surface (otherwise it would have a component along the surface that could do work, contradicting the first argument). The displacement of a charge along the surface lies in the surface itself, so the displacement is perpendicular to the field. The work done is $W = \vec{F} \cdot \vec{s} = Fs\cos\theta$, and with the angle between force and displacement equal to 90° the cosine is zero, so the work is zero. The two arguments give the same answer because they describe the same physical situation viewed from different angles.

9. The electric field component along a chosen direction equals the negative of the rate of change of the electric potential with distance in that direction. Equivalently, the electric field points in the direction in which the potential decreases most rapidly, and its strength equals the magnitude of that rate of decrease. The negative sign carries a physical meaning: a positive test charge released in an electric field naturally accelerates from

regions of higher potential toward regions of lower potential (as a mass rolls downhill on a gravitational slope). For the calculated field to point in this natural direction (from high V to low V), the rate of change dV/dr must be reversed in sign, since dV/dr is itself positive when V increases with distance. The negative sign is therefore the mathematical statement that fields point downhill, from high potential to low.

10. (i) The Earth is treated as the universal reference of zero potential because it is an effectively infinite reservoir of charge: any quantity of charge that flows to or from it produces a change in its own potential too small to measure. Like a vast lake compared with a teacup, the Earth's level does not rise or fall when charge is added or removed in ordinary quantities, so by convention this stable level is taken as the zero of potential, and the potential of any other body is reported relative to it.

(ii) At the moment a charged conductor is earthed (connected to the Earth by a conducting wire), the conductor and the Earth form a single conductor that must reach a single potential. If the conductor was at a positive potential, electrons flow from the Earth to the conductor (so the conventional current flows from conductor to Earth) and continue until the conductor is at the same potential as the Earth, that is, zero. If the conductor was at a negative potential, excess electrons flow from the conductor to the Earth (current flows from Earth to conductor) and continue until the conductor is again at zero potential. In both cases the conductor ends with no net charge.

EXERCISE 10B

11. The thundercloud above carries a large electric charge of one sign and induces a strong electric field in the air below it. A tall pointed object pushes a conducting path upward into this field, and at its sharp top the field is concentrated to values far higher than over the open ground around it. Air ionises and breaks down at a fixed field strength of about 3×10^6V/m, so it is at the top of the tallest object that this critical field is reached first. The discharge therefore begins at, and travels back down through, that tall pointed conductor. A shorter neighbouring tree never raises the field at its tip to the breakdown level, and is bypassed.

12. Kipanga is right that each hair, brought into rubbing contact with the wool jumper, has gained charge of the same sign as every other hair (one sign of charge has gone onto the hairs, and the opposite sign onto the wool). Like charges on the individual hairs repel one another, and so the hairs splay apart. Mr. Akilikubwa is adding that the body itself has acquired the same sign of charge through the rubbing and behaves as a charged conductor. Electric field lines around such a conductor point radially outward from its surface, and each charged hair, free to swing about its root, aligns itself along these outward field lines. It therefore stands up away from the head, not merely away from the neighbouring hairs.

13. The sharp upper point concentrates the electric field around the conductor to its highest value just above the tip, so this is the spot at which the surrounding air breaks down first. The point therefore takes any lightning strike in preference to the rest of the building, and, more importantly on every other day, it leaks small currents of charge quietly into the air through the same point-discharge mechanism, partially discharging the cloud above before any full strike can develop. The thick copper strap and the buried plate together provide a path of extremely low resistance leading any actual lightning current straight down through the strap to the earth, where it spreads harmlessly into the soil. Without the earth connection, the lightning current would force its way through the building's ordinary structure, with potentially fatal results.

14. A car body is essentially a hollow metal shell, a Faraday cage. When charge is delivered to its outer skin by a lightning strike, that charge moves entirely to the outer surface of the metal almost instantaneously, and the electric field inside the cavity of the shell is zero. The lightning current likewise spreads over the outside of the car and flows down to the ground through the tyres or through the air, never crossing the empty interior. The occupants, sitting in this field-free space, experience no electric force from the strike and are unhurt, provided they are not in contact with the metal frame at the moment the lightning is flowing.

15. The two charged plates set up an almost uniform electric field across the gap, and a charged droplet flying through that field experiences a steady sideways electric force equal to its charge multiplied by the field strength. While the droplet is between the plates it follows a parabolic path, just as a body in a uniform gravitational field would. Doubling the charge on the droplet doubles the sideways force and doubles the sideways deflection by the time the droplet leaves the field; halving the charge halves the deflection. By choosing the charge on each successive droplet, the print head selects exactly where on the page that droplet strikes, painting the shape of each letter one droplet at a time.

16. The very high voltage on the central wire ionises the air close to its surface, producing a stream of free electrons and gas ions. As dust and ash particles in the flue gas drift past this region they collide with the ions, and so each particle acquires a net electric charge of the same sign as that of the central wire. The strong electric field between the central wire and the surrounding earthed plates then exerts a Coulomb force on each charged particle, pulling it sideways across the gas stream and onto a collecting plate. The cleaned gas continues on up the chimney while the collected dust is shaken or rinsed from the plates and removed.

17. As fuel flows from the truck through the hose into the aircraft, friction between the moving fuel and the inner walls of the hose generates electric charge: the fuel acquires charge of one sign while the truck acquires the opposite. If the truck and the aircraft are not connected, a potential difference quickly grows between them as the truck loses charge and the aircraft gains it. Once that potential difference reaches the breakdown field of the air in the gap between the two metal bodies, a spark crosses the gap, and in the petrol-vapour atmosphere of a refuelling site that single spark is enough to start a fire. The bonding cable holds the truck and the aircraft at the same potential throughout the operation, so no potential difference ever develops between them and no spark can occur. The operator's rubber boots cannot perform this function, since they are insulators and would simply leave the two metal bodies free to drift to different potentials.

18. The lightning current entering the ground at the foot of the tree spreads radially outward through the soil, and the resistance of the soil makes the electric potential fall off steeply with distance from the strike point. Two points on the ground a small distance apart can therefore differ in potential by hundreds or thousands of volts during the brief flow of the strike. A person standing with feet a short distance apart has one foot at a higher potential than the other; current then flows from one foot, up one leg, across the body, down the other leg, and into the other foot. The current crossing the heart is what makes this step potential lethal, and explains why sheltering beneath an isolated tall tree during a thunderstorm is dangerous even when one is not struck directly.

19. Both the comb and the metal rod gain charge by friction with the hair just as readily; the question is what happens to that charge once it is on the surface. The plastic comb is an insulator, so any charge deposited on its surface stays in place, because electrons cannot flow through the plastic. The metal rod is a conductor in direct contact with the hand and so with the body of the experimenter, which is itself a fair conductor connected to the ground through the floor. Charge placed on the metal rod therefore flows along the rod, into the hand, and away through the body to the ground as quickly as the rubbing generates it. The rod is left uncharged a fraction of a second after each stroke, with nothing left with which to attract the paper.

20. The Van de Graaff dome at hundreds of kilovolts sets up a very strong electric field in the space around it. When the fluorescent tube is brought into that field, the field penetrates the thin glass and acts on the small amount of mercury vapour and noble gas inside, accelerating any free electrons that happen to be present. The accelerated electrons collide with the gas atoms and excite them, and as the atoms return to their lower energy states they emit ultraviolet light; the phosphor coating on the inside wall of the tube then converts this ultraviolet light into the visible glow that is seen. The glow does not require a current to flow through the tube from any external supply, only a strong electric field reaching into it from a charge in the neighbourhood.

EXERCISE 10C

21. The net field vanishes at the single point $x \approx 1.02\text{m}$, beyond the $-2.0\mu\text{C}$ charge on the side away from the origin.

22. By the cyclic ordering, the two +2.0μC charges occupy diagonally opposite corners of the square, and the two −3.0μC charges occupy the other diagonal pair. The field at the centre due to each charge lies along the diagonal joining that charge to the centre; the two contributions from the +2.0μC pair are therefore equal in magnitude and exactly antiparallel, and cancel, and the same is true of the two contributions from the −3.0μC pair. The net electric field at the centre of the square is therefore, $\mathbf{E = 0}$

23. $7.5 \times 10^4\ \text{N/C}$, 36.9° with the line BA extended past B, lying in the quadrant facing away from both A and C.

24. $E = \frac{4kQ}{3}$

25. (i) $5.65 \times 10^6\ \text{N/C}$ (ii) $56.5\mu\text{C}$ (iii) $6.39 \times 10^6\ \text{Nm}^2/\text{C}$

26. The student's formula $E = kQ/r^2$ treats the whole charge Q as if it were concentrated at the centre of the sphere. This is valid only for points outside the sphere, where the spherically symmetric distribution can be replaced by a single point charge at its centre. For a point inside the sphere, the field at radius r is determined only by the charge that lies inside that radius.

Apply Gauss's law to a spherical Gaussian surface of radius r centred on the sphere's centre. For a uniform volume charge density $\rho = Q/\left(\frac{4}{3}\pi R^3\right)$, the charge enclosed within the Gaussian surface is

$$Q_{enc} = \rho \times \frac{4}{3}\pi r^3 = Q\left(\frac{r}{R}\right)^3$$

so Gauss's law gives

$$E \times 4\pi r^2 = \frac{Q_{enc}}{\varepsilon_0} = \frac{Q}{\varepsilon_0}\left(\frac{r}{R}\right)^3$$

$$E = \frac{Qr}{4\pi\varepsilon_0 R^3} = \frac{kQr}{R^3}$$

The correct field grows linearly with r from zero at the centre, reaching kQ/R^2 at the surface (r = R), where it joins continuously with the external kQ/r^2 result. The student's formula gives the wrong dependence on r, and would diverge as r approaches zero, which is unphysical for a smooth volume distribution.

27. (i) $6.78 \times 10^5\ \text{N/C}$ (ii) 0

28. The total electrostatic potential energy of a system of point charges is the sum over all distinct pairs of kQ_iQ_j/r_{ij}. For the three vertices of an equilateral triangle of side L, every pair-separation equals L, and the three distinct pairs are (Q, +q), (Q, +q) and (+q, +q):

$$U_{total} = \frac{kQq}{L} + \frac{kQq}{L} + \frac{kq^2}{L} = \frac{k}{L}(2Qq + q^2)$$

Setting $U_{total} = 0$ and dividing through by kq/L (which is non-zero):

$$2Q + q = 0 \Rightarrow Q = -\frac{q}{2}$$

The unknown charge must be negative and half the magnitude of either of the other two, so that the attractive interactions of Q with the two +q charges exactly cancel the repulsive interaction between those two +q charges.

29. The total charge on the inner shell is $Q_1 = \sigma \times 4\pi r^2$ and on the outer shell is $Q_2 = \sigma \times 4\pi R^2$. For a thin spherical shell of charge Q and radius R, the potential at any point inside the shell (including at the centre) is the constant surface value kQ/R. By the superposition of potentials, the potential at the common centre is the sum of the contributions of the two shells, each evaluated at its own radius:

$$V_{centre} = \frac{kQ_1}{r} + \frac{kQ_2}{R} = \frac{k(\sigma 4\pi r^2)}{r} + \frac{k(\sigma\, 4\pi R^2)}{R}$$

$$\mathbf{V_{centre}} = 4\pi k\sigma(r + R) = \frac{\boldsymbol{\sigma(r+R)}}{\boldsymbol{\varepsilon_0}}$$

The contribution of each shell to the potential at the centre is proportional to its radius, not to its enclosed volume, because the field inside a uniformly charged shell is zero and the potential there is constant at its surface value.

30. (a) Let the displaced position of the test charge be x, with $|x| \ll a$. The +Q at x = −a lies to the left of the test charge at a distance $(a + x)$, and pushes the test charge in the +x direction; the +Q at x = +a lies to the right at a distance $(a - x)$ and pushes the test charge in the −x direction. Taking +x as positive, the net Coulomb force on the test charge is

$$F = \frac{kQq}{(a+x)^2} - \frac{kQq}{(a-x)^2}$$

For small displacement, expand each term to first order in x/a using the binomial series:

$$\frac{1}{(a \pm x)^2} \approx \frac{1}{a^2}\left(1 \mp \frac{2x}{a}\right)$$

so that

$$F \approx \frac{kQq}{a^2}\left(\left(1 - \frac{2x}{a}\right) - \left(1 + \frac{2x}{a}\right)\right) = -\frac{4kQq}{a^3}x$$

The net force is directed opposite to the displacement and proportional to it, which is the condition for simple harmonic motion.

(b) By Newton's second law:

$$m\frac{d^2x}{dt^2} = -\frac{4kQq}{a^3}x$$

Comparing with the standard SHM form $\frac{d^2x}{dt^2} = -\omega^2 x$:

$$\omega^2 = \frac{4kQq}{ma^3} \Rightarrow \boldsymbol{\omega = 2\sqrt{\frac{kQq}{ma^3}}}$$

31. The correct value of the work done is $W = 0$.

The student's error is to apply $W = qV$ as if V were the potential at a single point on the path. The relevant quantity is the potential difference between the start point A and the end point B of the path:

$$W = q(V_A - V_B)$$

For any closed loop, A and B are the same point, so $V_A = V_B$ and the potential difference $V_A - V_B = 0$. The work done is therefore zero, regardless of the value of the equipotential on which the loop happens to lie. The 12V value never enters the calculation.

The same conclusion follows from the conservative nature of the electrostatic force: the line integral of E around any closed loop in an electrostatic field is identically zero.

32. For the drop to remain at rest, the upward electric force on its negative charge must exactly balance its downward weight:

$$qE = mg$$

Between parallel plates the field is uniform, with $E = V/d$, so

$$q = \frac{mgd}{V} = \frac{(1.0 \times 10^{-14}\ \text{kg})(9.8\ \text{m/s}^2)(8.0 \times 10^{-3}\ \text{m})}{980\ \text{V}} = 8.0 \times 10^{-19}\ \text{C}$$

The number of excess electrons is therefore

$$n = \frac{q}{e} = \frac{8.0 \times 10^{-19}\ \text{C}}{1.6 \times 10^{-19}\ \text{C}} = 5$$

The drop carries 5 excess electrons. The fact that q comes out as an integer multiple of e is the experimental signature that charge is quantised, which is exactly what Millikan was looking for.

33. The small-angle limit gives the centre-to-centre separation d of the two balls at equilibrium in vacuum or air as

$$d^3 = \frac{Q^2 L}{2\pi\varepsilon_0 mg}$$

In a non-conducting medium of relative permittivity K the Coulomb force between the balls is reduced by a factor K, which is equivalent to replacing ε_0 by $K\varepsilon_0$, so

$$d_{oil}^3 = \frac{Q^2 L}{2\pi\, K\varepsilon_0 mg} = \frac{d_{air}^3}{K}$$

For small angle θ measured from the vertical, $d = 2L\sin\theta \approx 2L\theta$, so the ratio of separations is the ratio of half-angles:

$$K = \left(\frac{d_{air}}{d_{oil}}\right)^3 = \left(\frac{\theta_{air}}{\theta_{oil}}\right)^3 = \left(\frac{15°}{9°}\right)^3 = \frac{125}{27} \approx 4.6$$

The relative permittivity of the oil is approximately 4.6.

34. Assumptions: every ash particle reaches terminal drift velocity almost immediately (the time to reach terminal velocity is negligible compared with 2.0s); buoyancy and gravity on the micrometre particle are small compared with the electric driving force; the particle remains a uniform sphere of radius $r = 0.5\mu m$ throughout its drift; the applied field is uniform between the wire and the plate.

At terminal velocity the electric force balances the Stokes drag:

$$qE = 6\pi\eta\, rv \Rightarrow v = \frac{qE}{6\pi\eta\, r}$$

To cross the 15cm gap during the 2.0s residence time, the drift speed must be

$$v_{min} = \frac{0.15\ \text{m}}{2.0\ \text{s}} = 0.075\ \text{m/s}$$

so the minimum applied field is

$$E_{min} = \frac{6\pi\eta\, rv_{min}}{q} = \frac{6\pi\,(1.8 \times 10^{-5})(0.5 \times 10^{-6})(0.075)}{(100)(1.6 \times 10^{-19})} \approx 8.0 \times 10^5\ \text{V/m}$$

The required field of about 8×10^5V/m is roughly one quarter of the breakdown field of dry air (3×10^6V/m), so it lies safely below the breakdown limit. The chosen field is therefore realistic for a precipitator operating in ordinary flue gas conditions.

35. For an isolated spherical conductor of radius R at a potential V above ground, the field at the surface is $E = V/R$. With R = 0.06m and V = 350kV:

$$E_{surface} = \frac{V}{R} = \frac{350 \times 10^3\ \text{V}}{0.060\ \text{m}} = 5.83 \times 10^6\ \text{V/m}$$

This is about 1.9 times the dielectric strength of dry air (3.0×10^6V/m). The air immediately above the dome will therefore ionise and a continuous corona discharge (or sparking to any nearby earthed object) will drain charge from the dome as fast as the belt delivers it. The dome cannot reach, and cannot hold, the rated 350kV in ordinary dry air.

Either of the following design modifications will allow the rating to be met:

(i) Increase the radius of the dome. The minimum radius required to keep $E_{surface}$ below the breakdown field is $R_{min} = V/E_{breakdown} = 350 \times 10^3/3.0 \times 10^6 = 0.117$m, so a dome of diameter at least 24cm (twice the present 12cm) would hold 350kV safely.

(ii) Enclose the generator in a chamber of higher-dielectric-strength gas (for example sulphur hexafluoride, SF_6, whose breakdown field at atmospheric pressure is about three times that of air), or operate the generator at elevated pressure, since the breakdown field rises with gas pressure.

EXERCISE 11A

1. On a charged capacitor, the two plates carry charges of equal magnitude but opposite sign: one plate holds $+Q$ and the other plate holds $-Q$, so the net charge on the device as a whole is $+Q + (-Q) = 0$. The phrase "the capacitor stores charge" refers not to a net charge on the device but to the magnitude Q held separately on each plate, since it is the separation of equal and opposite charges across the gap, and not any net charge, that the capacitor is designed to maintain. The energy stored is contained in the electric field between the plates, set up by the equal and opposite surface-charge densities on the two facing plates. The two statements are therefore consistent: the device stores $+Q$ on one plate and $-Q$ on the other, with zero net charge as a whole, but with energy and a field between.

2. The capacitance of a parallel-plate capacitor is $C = \varepsilon_r\varepsilon_0 A/d$, so four factors determine it:

(i) the relative permittivity ε_r of the dielectric, with C directly proportional to ε_r;

(ii) the area A of overlap of the two plates, with C directly proportional to A;

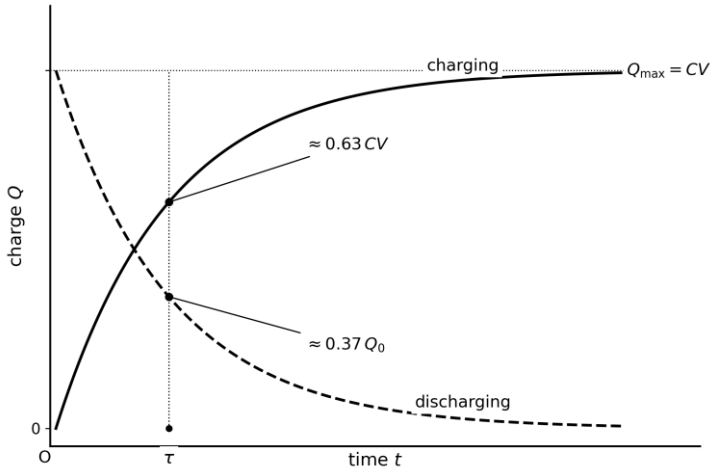
(iii) the separation d between the plates, with C inversely proportional to d; and

(iv) the permittivity of free space ε_0, a universal constant entering as a multiplicative factor.

Note: *The capacitance does* ***not*** *depend on the charge Q on the plates or the potential difference V across them, because $C = Q/V$ is fixed by the geometry and the dielectric alone and is the same whatever values Q and V take.*

3. In the parallel configuration each capacitor sees the full battery voltage V, so the total charge delivered by the battery is $Q_p = (C_1 + C_2)V$. In the series configuration the equivalent capacitance is $C_s = C_1C_2/(C_1 + C_2)$, and the battery delivers $Q_s = C_sV$ to the chain. Since $C_1 + C_2 > C_1C_2/(C_1 + C_2)$ for any positive C_1 and C_2, the **paralle**l configuration draws the **greater** total charge from the battery at the same voltage.

4. The charging curve follows $Q(t) = CV(1 - e^{-t/\tau})$, rising from zero and approaching the asymptotic value $Q_{max} = CV$ as t becomes large. The discharging curve follows $Q(t) = Q_0\, e^{-t/\tau}$, falling from Q_0 and decaying to zero. Both curves share the same time constant $\tau = RC$. At $t = \tau$ the charging curve has reached approximately $0.63\,CV$ (that is, $CV(1 - 1/e)$), and the discharging curve has fallen to approximately $0.37\,Q_0$ (that is, Q_0/e).



5. Mica is preferred for two reasons. First, mica has a dielectric strength of about $100\,MV/m$, several times larger than the value of about $16\,MV/m$ for impregnated paper, so a mica capacitor can hold off the peak RF voltage in a tuned circuit without breakdown, whereas a paper capacitor of the same gap would arc over. Second, mica has very low dielectric loss at radio frequencies, while paper's dielectric loss rises sharply with frequency; a paper capacitor in an RF tuning circuit would dissipate a significant fraction of the signal energy as heat, broadening the resonance peak and degrading the Q-factor of the tuned circuit.

6. An electrostatic precipitator is built as a thin central wire held at very high electric potential (typically several tens of kilovolts negative) inside an enclosure of earthed collector plates. The intense electric field close to the surface of the wire ionises the gas locally, producing a stream of free electrons that attach themselves to passing smoke particles and charge each particle up. The strong field between the wire and the collector plates then exerts a Coulomb force on each charged particle, driving it sideways out of the gas stream and onto a plate, where it sticks and is later shaken off mechanically. The cleaned gas continues to the chimney while the collected dust is removed from the plates and disposed of.

7. The dielectric constant, or relative permittivity ε_r, is the dimensionless ratio of the capacitance with the given material filling the gap to the capacitance of the same capacitor in vacuum, and measures how strongly the material polarises in the applied field. The dielectric strength is the maximum electric field, in V/m, that the material can sustain before electrical breakdown turns it from an insulator into a conductor.

Pure water has an exceptionally high ε_r of about 81 at room temperature, but its dielectric strength is very low because even trace amounts of dissolved ions allow it to conduct electrolytically as soon as a modest field is applied. A water-filled capacitor would therefore self-discharge through ionic conduction long before any useful charge could be stored, so the high ε_r is of no practical benefit.

8. From $V = IR$ we have $R = \frac{V}{I}$. From $Q = CV$ we have $C = \frac{Q}{V}$,.

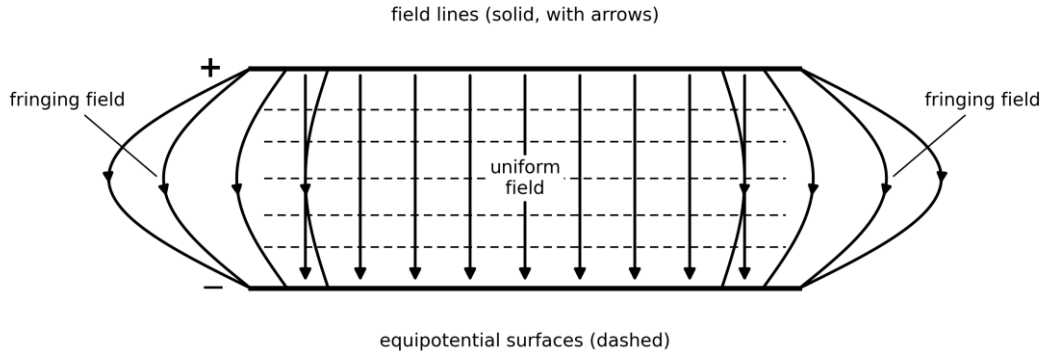
$$[RC] = [R] \times [C] = \frac{[V]}{[i]} \times \frac{[Q]}{[V]} = \frac{[Q]}{[i]} = \frac{[I][T]}{[I]} = [T]$$

Since RC has dimensions of time, it must also have units of time.

9. During charging, the growing potential difference Q/C across the capacitor opposes the EMF V_o of the source, so the net driving voltage in the loop is $V_o - Q/C$ and the current in the resistor is $I(t) = (V_o - Q/C)/R$. As Q approaches its asymptotic maximum CV_o, the numerator $V_o - Q/C$ smoothly approaches zero, and so does the current. The current cannot snap off because the approach of Q to CV_o is exponential and asymptotic, not a sudden hit against a ceiling: the same exponential law that drove the early rapid charging governs the slow tail at the end, and the current dies away smoothly as $e^{-t/RC}$.

10. The steady-state charge is determined by the capacitance C and the EMF V_o of the source together, as $Q_{max} = CV_o$. The series resistance R does not appear in this expression because, at the end of charging, the current has fallen to zero and no voltage is dropped across R; the entire EMF then sits across the capacitor, requiring $Q = CV_o$ regardless of R. What R does control is the time constant $\tau = RC$, which sets how fast Q_{max} is approached. A larger R simply means a longer time constant; the capacitor still reaches the same Q_{max} eventually, only more slowly.

11. In the central region between the plates, away from the edges, the field lines run straight from the + plate to the − plate, parallel to one another and uniformly spaced, and the field there is uniform. The equipotential surfaces in this central region are flat planes parallel to the plates, intersecting the field lines at right angles. Near the edges of the plates the field lines bulge outwards beyond the plate boundaries (the fringing field), and the equipotential surfaces curve outwards to follow them, while remaining perpendicular to the field at every point.



12. Yes, the two plates carry charges of equal magnitude after being connected to the battery, despite their unequal areas. Charge is conserved in the isolated capacitor-battery loop: any positive charge that flows from one terminal of the battery onto one plate must be matched by the same amount of positive charge flowing from the other plate into the other terminal of the battery, so the two plates finish with charges $+Q$ and $-Q$ regardless of the plate areas. The surface charge densities $\sigma_1 = Q/A_1$ and $\sigma_2 = Q/A_2$ differ because the areas A_1 and A_2 differ, with σ larger on the smaller plate; the totals on the two plates are nevertheless equal in magnitude and opposite in sign, as conservation of charge demands.

EXERCISE 11B

13. A camera flash is a capacitor charged through the camera's circuit by the batteries. The total resistance of the charging loop is the internal resistance r of the batteries plus the fixed circuit resistance R_{circ}, so $R_{tot} = r + R_{circ}$. As the cells weaken, their internal resistance r grows steadily, so R_{tot} grows and the time constant $\tau = R_{tot}\,C$ grows with it. The capacitor still reaches its full charge eventually, because $Q_{max} = CV$ is set by C and the EMF, not by R; but it does so more slowly. The owner therefore notices a longer wait between flashes as the cells weaken, since the flash circuit cannot fire until the capacitor has reached close to its full charge.

14. The time constant for a series RC circuit is $\tau = R\,C$, so the required capacitance is:

$$C = \frac{\tau}{R} = \frac{10\text{s}}{100 \times 10^3\Omega} = 1.0 \times 10^{-4}\text{F} = 100\mu\text{F}$$

The required capacitance is 100μF.

15. Rotating the moving plates of the variable capacitor partly out of the fixed plates reduces the effective area of overlap A between the two sets of plates, and the capacitance $C = \varepsilon_0 A/d$ therefore decreases. The radio's resonant frequency is $f_0 = 1/\left(2\pi\sqrt{L\,C}\right)$, which is inversely proportional to $\sqrt{C}$. As C decreases, f_0 rises, and the radio tunes from a station of lower carrier frequency to one of higher carrier frequency.

16. When the trigger is pressed, the sudden compression of the piezoelectric element produces a brief charge pulse on its faces; the element behaves at that instant as a small capacitor whose plates have been charged from rest to a high voltage. The stored charge dumps almost instantly through the spark gap as a single high-current discharge, producing the spark. Once the element has discharged through the gap, it carries no further charge while the trigger is held down; no current flows and no further spark is produced. Only when the trigger is released and re-pressed is the element compressed again, generating a fresh charge pulse and a fresh single spark.

17. An electrolytic capacitor's dielectric is a thin oxide layer formed and maintained on the surface of one of the aluminium foils by the chemistry of the electrolyte under the correct polarity. The oxide layer is what gives the device its very high capacitance per unit volume. Reverse polarity reverses the chemistry: the oxide dissolves rather than forms, the dielectric layer rapidly thins until it can no longer block conduction, and the capacitor then behaves as a short circuit at the applied voltage. A large current flows, boils the electrolyte, and can rupture the can, sometimes violently. The polarity marking is therefore a safety warning, not merely an electrical convenience.

18. Power dissipated by the bulb:

$$P = \frac{V^2}{R} = \frac{(3\text{V})^2}{100\Omega} = 0.090\text{W} = 90\text{mW}$$

Energy required for 30s of glow:

$$U = P \times t = (0.090\text{W})(30\text{s}) = 2.7\text{J}$$

The capacitor must store at least this much energy. From $U = 1/2\ CV^2$:

$$C = \frac{2\,U}{V^2} = \frac{2 \times 2.7\text{J}}{(3\text{V})^2} = 0.6\text{F}$$

The capacitance is about 0.6F.

19. Water has a relative permittivity ε_r of about 81 at room temperature, far larger than that of dry soil components (ε_r about 4). When water fills the spaces between soil particles in the sensor's gap, the effective ε_r of the soil-water mixture rises sharply, from about 4 in dry conditions to about 15 to 20 when saturated. The capacitance $C = \varepsilon_r\varepsilon_0 A/d$ scales linearly with this effective ε_r, so a factor-of-four increase in capacitance (from 4nF to 15nF) reflects a factor-of-four increase in the effective dielectric constant, exactly the change expected when dry soil is wetted to saturation by the previous day's rain.

20. With the switch open, the battery is disconnected from the rest of the circuit, and the only available path for the capacitor to discharge is through the bulb itself. The capacitor therefore discharges as a simple RC circuit with R equal to the bulb's resistance, the current decaying exponentially as $I(t) = (V_0/R)\,e^{-t/R\,C}$. The bulb is lit by this decaying current and dims smoothly as the current falls. The afterglow lasts of the order of one or two time constants $\tau = R\,C$; for a torch designed with τ of about half a second, the bulb dims noticeably within a second and then goes dark.

21. The two plates of the giant capacitor are the lower charge layer of the storm cloud, typically carrying a strong negative charge at its base, and the surface of the Earth below, which is induced positively by the cloud's charge. The dielectric is the column of air between them. Air, like any insulator, has a dielectric strength: a maximum electric field above which the air molecules ionise and the gas conducts as a plasma. For dry air this critical field is approximately $3 \times 10^6\text{V/m}$. As long as the field between cloud and ground remains below this value, the air holds off the

stored charge and the cloud-Earth system behaves as a charged capacitor; once the field exceeds the dielectric strength of air, breakdown occurs along a conducting channel, and the stored energy releases as the lightning stroke. Humid air has a lower breakdown field, which is part of the reason why lightning is most common in stormy, humid conditions.

EXERCISE 11C

22. Series: 1.09μF, 262μC, 31.4mJ. Parallel: 12μF, 2.88mC, 0.346J.
23. (i) 2.67μF (ii) 320μC (iii) 80V (iv) 19.2mJ
24. 7.37s.
25. Half-life:138.6s, 7 half-lives are required.
26. 3.6×10^{-4}N, the plate force is about **27** times **smaller** than the weight of 1g of water.
27. (i) 8.19mC (ii) 8.2J
28. The stack is two capacitors in series, the upper of capacitance $C_1 = 2K_1\varepsilon_0 A/d$ and the lower $C_2 = 2K_2\varepsilon_0 A/d$, each with full plate area A and thickness $d/2$. In series:

$$\frac{1}{C_{eq}} = \frac{1}{C_1} + \frac{1}{C_2} = \frac{d}{2\varepsilon_0 A}\left(\frac{1}{K_1} + \frac{1}{K_2}\right) = \frac{d}{2\varepsilon_0 A}\left(\frac{K_1 + K_2}{K_1 K_2}\right)$$

Hence:

$$C_{eq} = \frac{2\varepsilon_0 A}{d}\left(\frac{K_1 K_2}{K_1 + K_2}\right)$$

29. (i) $I_{peak} = 250A$ (ii) $\tau = 40\mu s$ (iii) The entire stored energy is dissipated as heat in the coil: 2.5J.
30. Earthing the outer shell forces its potential to zero, and the inner shell at potential V_{in} relative to it has:

$$V_{in} = \frac{kQ}{a} - \frac{kQ}{b} = kQ\left(\frac{b-a}{ab}\right)$$

So $C = \frac{Q}{V_{in}} = \frac{ab}{k(b-a)} = \frac{4\pi\varepsilon_0 ab}{b-a}$, confirming a capacitor relationship $Q = CV$.
With $a = 0.05$m, $b = 0.10$m:

$$C = \frac{0.05 \times 0.10}{(9 \times 10^9)(0.05)} = 11.1\text{pF}$$

An isolated sphere of radius $a = 0.05$m has $C_{iso} = 4\pi\varepsilon_0 a = a/k = 5.56$pF. The double-shell capacitor has twice the capacitance of the isolated inner sphere alone.
31. 6.9nF.
32. 96pF.
33. 5cm
34. 9.5×10^{-3}K
35. (i) 11.1pF (ii) 44.3pF (iii) 200pF (iv) Place the sensor as the timing capacitor of an RC or LC oscillator. As soil moisture rises, the sensor's capacitance rises and the oscillator frequency falls. A frequency-to-voltage converter (a simple monostable triggered by the oscillator, with the output averaged through a low-pass filter) then drives a moving-coil meter that has been calibrated against soil samples of known moisture content, giving the farmer a direct percentage reading.
36. (i) 0.721s (ii) 1.50s
37. (i) 8.85nC (ii) 112V (iii) 0.495μJ
38. (i) 109.6mJ (ii) 96.1mJ (iii) 13.5mJ. The energy lost is dissipated as heat in the connecting wires during the brief transient redistribution of charge between the two capacitors.

EXERCISE 12A

1. The electrons are already present all along the wire before the switch is closed, like a queue of people filling a corridor. When the switch closes, the electric field is set up along the whole wire almost at the speed of light, so every electron begins to drift at once and the lamp lights immediately. The electrons themselves only crawl forward at about a millimetre per second; it is the field, not the electrons, that travels fast, just as the bucket is passed down a standing queue at once although each person barely moves.
2. The EMF is the energy supplied per coulomb by the source, equal to the work done per unit charge when no current flows. The terminal voltage is the potential difference actually available at the terminals when current flows, and is smaller than the EMF by the internal drop Ir. The two are equal only when no current is drawn, that is on open circuit, because then $Ir = 0$.
3. Ohm's law states that the current through a conductor is directly proportional to the potential difference across it, provided the temperature and other physical conditions remain constant, so that $V = IR$ with R constant. A metal wire at constant temperature and constantan obey it; a filament lamp and a semiconductor diode do not. It is called empirical because it is an experimental result that holds only for certain materials under certain conditions, rather than a universal law derived from first principles.
4. The conductance of a conductor measures how easily it allows current to pass, and is the reciprocal of its resistance, $G = \frac{1}{R}$. Its SI unit is the siemens, equal to one ampere per volt. It is seldom used in school problems because series resistances simply add and the familiar form $V = IR$ is more direct, so conductance offers no real advantage there.
5. In a metal the number of free electrons is essentially fixed, so the resistance depends only on how often the electrons are scattered. Raising the temperature makes the lattice ions vibrate more violently, so the electrons collide with them more often and the relaxation time τ falls. Since $R = \frac{ml}{ne^2A\tau}$, a smaller τ gives a larger R, so the resistance of a metal rises with temperature.
6. Resistance is the opposition that a particular specimen offers to current; it depends on the material and on the dimensions of the specimen, and its unit is the ohm. Resistivity is a property of the material alone, independent of shape or size, numerically equal to the resistance of a unit cube of the material, and its unit is the ohm-metre.
7. Nichrome and constantan have a resistivity about sixty times that of copper, so a short length gives a usefully large resistance, whereas copper would need an impossibly long wire. They also have a very small temperature coefficient of resistance, so their resistance barely changes as they heat, keeping a heater's power or a standard resistor's value steady. Copper fails on both counts, having a low resistivity and a large temperature coefficient.

8. Kirchhoff's current law states that the total current entering any junction equals the total current leaving it, and it embodies the conservation of electric charge. Kirchhoff's voltage law states that around any closed loop the sum of the EMFs equals the sum of the potential drops, and it embodies the conservation of energy.

9. An ordinary ohmmeter passes a current through the unknown resistor and reads it on a calibrated scale, so its accuracy is limited by that calibration and by the current drawn. A Wheatstone bridge is instead adjusted until the galvanometer reads exactly zero, and at balance no current flows through the galvanometer, so the unknown resistance is fixed by a ratio of the other three resistances and does not depend on the galvanometer's calibration at all, only on its ability to detect zero. This null method is what makes the bridge more accurate.

10. A potentiometer assumes that the potential falls evenly along its wire, so that the balance length is directly proportional to the potential difference being measured. If the gradient were not uniform, equal lengths would not correspond to equal voltages and the readings would be meaningless. A uniform gradient is guaranteed when the wire is of uniform cross-section and material, so its resistance per unit length is constant, and the driving current is kept steady.

11. A fuse is a short length of low-melting-point alloy connected in series with the circuit. The current through it heats it at the rate $P = I^2R$. At normal currents this heating is mild and the fuse survives, but if a fault makes the current rise too high, the increased I^2R heating melts the fuse, so the circuit breaks and the current stops before the wiring or appliance can overheat and catch fire.

12. The maximum power transfer theorem states that *a source of fixed EMF and internal resistance r delivers the greatest power to a load when the load resistance equals the internal resistance, $R = r$.* At that matched point the efficiency is only 50%, because half the power is wasted inside the source. A transmission grid must deliver almost all its power to consumers, so it deliberately makes the load resistance far larger than the line resistance and runs at high efficiency, well away from the matched condition; matching would waste half the generated power as heat in the lines.

13. In an initiated discharge the gas conducts only while an external ionising agent keeps producing ions, so the current stops the moment the agent is removed, since the applied voltage merely collects ready-made ions. In a self-sustaining discharge the voltage is large enough that the discharge makes its own ions and continues without any external agent. The mechanism is chain ionisation: each accelerated electron gains enough energy between collisions to ionise the next molecule, those freed electrons ionise still more, and the resulting avalanche keeps the gas supplied with carriers.

14. The colour a gas emits is set by the discrete energy levels of its atoms, because when an excited electron falls between two levels it emits a photon whose energy, and hence colour, equals the gap between them. Sodium atoms have an energy gap corresponding to yellow-orange light, while neon atoms have gaps corresponding to red light. Since each element has its own unique set of energy levels, each glows with its own characteristic colour.

EXERCISE 12B

15. The earth pin joins the metal casing to the ground through a path of very low resistance, and in normal use the casing carries no current. Should a fault let the live wire touch the casing, the casing would otherwise rise to the live potential and shock anyone who touched it. Because the earth connection offers an easy path to ground, a large fault current instead flows through it and blows the fuse, cutting off the supply. The casing is therefore never left at a dangerous potential.

16. The power delivered is $P = VI$, so for a fixed power a higher transmission voltage means a proportionally smaller current. The power wasted as heat in the line is I^2R_{line}, which grows with the square of the current. Raising the voltage therefore lowers the current and cuts the line loss sharply, so that far more of the generated power reaches the consumer. Transmitting at the household voltage would demand an enormous current and waste most of the power as heat in the lines.

17. The resistance of a metal filament rises with temperature, so a cold filament has its lowest resistance. At the instant of switching on the filament is still cold, so it offers little resistance and a large surge of current passes through it. This sudden surge produces a burst of heating that can snap the already thin and weakened filament. Once the filament warms, its resistance rises and limits the current, so failure is far more likely at the cold moment of switch-on than during steady glowing.

18. The brightness depends on the terminal voltage supplied to the bulb, given by $V = E - Ir$. As the cells are used, chemical changes within them increase the internal resistance r, so the internal drop Ir grows and the terminal voltage falls. Less power therefore reaches the bulb and it glows more dimly. New cells have a small internal resistance and deliver almost the full EMF to the bulb, so they make it brighter.

19. A lamp's filament has a roughly fixed resistance, and the power it dissipates is $P = \frac{V^2}{R}$, which rises with the square of the applied voltage. Connecting it to a much higher voltage therefore drives its power up steeply, so the filament is heated far beyond its design temperature and melts almost at once. The lamp is rated for one particular voltage precisely because a higher one destroys it.

20. The extension lead has its own resistance, and the heavy current drawn by the heater heats the lead itself at the rate I^2R. When the lead is coiled on its reel, the turns trap this heat and prevent it from escaping, so the temperature of the insulation climbs steadily. The insulation may then soften, melt, and short, starting a fire. The same lead run out straight sheds its heat easily and stays cool, which is why a heavy load must never be drawn through a coiled lead.

21. To measure resistance the meter passes a small current through the coil, and this current gently heats the copper. Since the resistance of copper rises with temperature, the coil's resistance increases as it warms. The longer the meter is left connected, the warmer the coil grows, so the displayed reading slowly climbs.

22. The counter operates in the self-sustaining region of the ionisation curve, where each ionising particle that enters the tube triggers a full avalanche and so produces one current pulse, heard as a click. Bringing the counter closer to the source means more particles reach the tube each second, so more avalanches are triggered per second and the clicks come more rapidly. The count rate therefore measures how many particles arrive, not how energetic each one is.

23. The light given off by a hot body shifts towards shorter, bluer wavelengths as its temperature rises. The centre of the arc is the hottest region, so it emits the high-energy bluish-white light, while towards the edges the arc is cooler and emits longer-wavelength orange light. The change of colour across the arc therefore maps the change of temperature across it.

24. A cold tube contains few free carriers and has a high breakdown potential, so the mains voltage alone cannot strike the discharge. Heating the cathode filaments makes them release electrons by thermionic emission, seeding the gas with free carriers. With these seed electrons present, the high-voltage pulse from the choke can then ionise a path along the tube and start the discharge. The preheating therefore supplies the carriers the tube needs before it can light.

25. The current driven through a person is set by the voltage and the resistance of the body, $I = \frac{V}{R}$. Dry skin has a high resistance, which limits the current to a small and usually harmless value. Wet skin, however, has a much lower resistance, so for the same voltage a far larger current flows through the body, and the shock is correspondingly more severe.

EXERCISE 12C

26. The resistance of the parallel combination is 1.33Ω.
27. The equivalent resistance between A and B is 8Ω.
28. The resistance is 20% greater at a temperature of 50°C.
29. The cell has an EMF of 5.0V and an internal resistance of 1.0Ω.
30. The equivalent resistance is 2Ω, and the currents are 6A, 4A and 2A.
31. The cells supply 1.0A and 2.0A, the central resistor carries 3.0A, and it dissipates 18W.
32. The unknown resistance is $S = 100\Omega$.
33. The series arrangement gives 1.5A and the parallel arrangement 0.71A.
34. The battery delivers 2.4A, which divides into 0.4A, 0.8A and 1.2A through the 12Ω, 6Ω and 4Ω resistors respectively.
35. The drift velocity of the electrons is $1.7 \times 10^{-4}\text{ms}^{-1}$.
36. The current is 2.0A, and the powers are 8W, 12W and 20W.
37. The equivalent resistance between X and Y is 2Ω.
38. The two resistances are 4Ω and 6Ω.
39. When the cells assist each other, their EMFs add and so do their internal resistances, giving a total EMF of $4\text{V} + 2\text{V} = 6\text{V}$ against a total resistance of $0.5\Omega + 0.5\Omega + 3\Omega = 4\Omega$:

$$\text{I}_{\text{aiding}} = \frac{6\text{V}}{4\Omega} = 1.5\text{A}$$

When one cell is reversed, the EMFs subtract while the resistances still add, giving a net EMF of $4\text{V} - 2\text{V} = 2\text{V}$ against the same 4Ω:

$$\text{I}_{\text{opposing}} = \frac{2\text{V}}{4\Omega} = 0.5\text{A}$$

The current is 1.5A when the cells assist each other and 0.5A when they oppose.
40. The two resistors in parallel have a combined resistance of:

$$\text{R}_\text{p} = \frac{(6\Omega)(3\Omega)}{6\Omega + 3\Omega} = 2\Omega$$

The total power dissipated is $\text{P} = \frac{\text{V}^2}{\text{R}_\text{p}}$, so the supply voltage is:

$$\text{V} = \sqrt{\text{PR}_\text{p}} = \sqrt{(18\text{W})(2\Omega)} = 6.0\text{V}$$

This voltage acts across each resistor, so the branch currents are:

$$\text{I}_6 = \frac{6.0\text{V}}{6\Omega} = 1.0\text{A}, \qquad \text{I}_3 = \frac{6.0\text{V}}{3\Omega} = 2.0\text{A}$$

The supply voltage is 6.0V, and the currents are 1.0A in the 6Ω resistor and 2.0A in the 3Ω resistor.
41. With nothing connected across PQ, no current is drawn from the tap, so the same current flows through both divider resistors, and the output is the fraction of the supply that falls across the lower resistor:

$$\text{V}_\text{out} = \text{V} \times \frac{\text{R}_2}{\text{R}_1 + \text{R}_2} = 12\text{V} \times \frac{2\text{k}\Omega}{4\text{k}\Omega + 2\text{k}\Omega} = 4.0\text{V}$$

When a 2kΩ load is connected across PQ, it is in parallel with the lower 2kΩ resistor, and the two combine to:

$$\text{R}_\text{lower} = \frac{(2\text{k}\Omega)(2\text{k}\Omega)}{2\text{k}\Omega + 2\text{k}\Omega} = 1\text{k}\Omega$$

The output now falls across this smaller resistance:

$$\text{V}_\text{out} = 12\text{V} \times \frac{1\text{k}\Omega}{4\text{k}\Omega + 1\text{k}\Omega} = 2.4\text{V}$$

The output is 4.0V on open circuit but drops to 2.4V under load, which shows that drawing current from a potential divider lowers its output voltage.
42. The average current is the charge transferred divided by the time it takes:

$$\text{I} = \frac{\text{Q}}{\text{t}} = \frac{15\text{C}}{0.20 \times 10^{-3}\text{s}} = 7.5 \times 10^4\text{A}$$

The energy released is the charge multiplied by the potential difference through which it falls:

$$\text{E} = \text{QV} = (15\text{C})(3.0 \times 10^8\text{V}) = 4.5 \times 10^9\text{J}$$

The average power is the energy delivered divided by the time:

$$\text{P} = \frac{\text{E}}{\text{t}} = \frac{4.5 \times 10^9\text{J}}{0.20 \times 10^{-3}\text{s}} = 2.25 \times 10^{13}\text{W}$$

The stroke carries an average current of 7.5×10^4A, releases 4.5×10^9J of energy, and delivers it at an average power of 2.25×10^{13}W, which is why lightning is so destructive in the brief instant it lasts.
43. The cell drives 2.0A, holds a terminal voltage of 11V, delivers 22W to the external circuit, and wastes 2W within itself.

EXERCISE 13A

1. In one complete cycle an alternating current flows first one way, falls to zero, then flows the other way and returns to zero, so it changes direction twice in every cycle. At a frequency of 50Hz the current completes 50 whole cycles in each second, and since every cycle contains two reversals the number of reversals each second is $2 \times 50 = 100$. The current therefore reverses one hundred times a second even though only 50 complete cycles, and hence a frequency of 50Hz, occur in that time.
2. The figure 240V is the rms value, the steady value that would give the same heating; it is not the greatest voltage the supply reaches. The voltage actually climbs to a peak of $\text{V}_0 = \text{V}_\text{rms}\sqrt{2} = 240\text{V} \times \sqrt{2} \approx 339\text{V}$ at the crest of every cycle. Since the insulation has to hold back the largest voltage that ever appears across it, it must be able to withstand this peak of about 339V and not merely the quoted rms value of 240V.
3. When the current in an inductor changes it sets up a changing magnetic flux, which induces a back-EMF that opposes the change in the current. This back-EMF is largest when the current is changing fastest, which happens as the current passes through zero, and the applied voltage must rise to its own peak to overcome it. The voltage therefore peaks while the current is still climbing through zero, so the current does not reach its own peak until a quarter of a cycle later; that is, the current lags the voltage by 90°.

4. An ordinary ohmmeter measures opposition by sending a steady direct current through the component and noting the result. A capacitor, however, blocks direct current as soon as it is fully charged, so no steady current flows and the meter simply reads an extremely high, effectively infinite, value. The true opposition of a capacitor is its reactance $X_C = \frac{1}{2\pi fC}$, which appears only with alternating current and depends on the frequency, so it cannot be obtained from the direct-current test that an ohmmeter carries out.

5. The voltage across the resistor is in phase with the current, while the voltage across the inductor reaches its peak a quarter of a cycle earlier, so the two voltages never reach their maximum values at the same instant. Because they are 90° out of phase, they must be combined as the two perpendicular sides of a right-angled triangle, giving an impedance equal to the hypotenuse, $Z = \sqrt{R^2 + X_L^2} = \sqrt{(30\Omega)^2 + (40\Omega)^2} = 50\Omega$. The combined opposition is therefore 50Ω, less than the arithmetic sum of 70Ω, because a hypotenuse is always shorter than the sum of the other two sides.

6. The impedance of the circuit is $Z = \sqrt{R^2 + (X_L - X_C)^2}$, in which the inductive reactance rises with frequency while the capacitive reactance falls. At one particular frequency the two reactances become equal and cancel, leaving the impedance at its smallest value, the resistance R alone, so the current $I = \frac{V}{Z}$ rises to a maximum. At any higher or lower frequency the reactances no longer cancel, the net reactance $(X_L - X_C)$ is no longer zero, the impedance is larger, and the current is therefore smaller.

7. A high quality factor makes the resonance curve tall and narrow, so the receiver responds strongly to the wanted station and only weakly to the stations on either side, giving the sharp tuning needed to pick one station out of many. If the quality factor were made far too high, the band of frequencies the circuit accepts would become so narrow that part of the broadcast signal would be cut off, so the reception would be weakened and distorted. A good receiver therefore needs a quality factor high enough to separate the stations but not so high as to spoil the signal it selects.

8. At the resonant frequency the inductor and the capacitor have equal reactances, so the current in the inductive branch and the current in the capacitive branch are equal in size but, being half a cycle apart in phase, opposite in direction at every instant. Where the branches rejoin these two currents very nearly cancel, so only a tiny net current is taken from the supply, even though each branch alone carries a large current that simply circulates to and fro around the loop. The parallel combination therefore behaves as a very high impedance at resonance, which is why it is used to reject that frequency.

9. The voltmeter reads the rms voltage and the ammeter reads the rms current, and their product, $V_{rms}I_{rms} = 240V \times 5.0A = 1200VA$, is only the apparent power. The power actually delivered is the true power $P = V_{rms}I_{rms}\cos\phi$, where ϕ is the phase angle between the current and the voltage, because only the part of the current in phase with the voltage does useful work. The delivered power equals 1200W only if the circuit is purely resistive, so that $\phi = 0$ and $\cos\phi = 1$; for any other load the phase difference makes it less.

10. A generator or transformer is limited by the greatest voltage its insulation can stand and the greatest current its windings can carry without overheating, and the product of these two is the apparent power, measured in volt-amperes. The useful power it actually supplies is $VI\cos\phi$, which depends on the power factor of whatever load happens to be connected, something the maker cannot know beforehand. The rating is therefore quoted in kilovolt-amperes, since that states the safe limits of voltage and current whatever the power factor of the load may be.

11. A transformer works by mutual induction, in which the changing flux set up by the primary must link the turns of the secondary if it is to induce an EMF in them. A common iron core offers an easy path that guides and concentrates almost the whole of this changing flux through the secondary coil, so the two coils are tightly coupled. If the coils were merely placed side by side in air, most of the primary's flux would spread out and miss the secondary, so very little EMF would be induced and the transformer would be extremely inefficient.

12. A transformer passes power from the primary to the secondary, and for an ideal transformer the power drawn from the mains equals the power delivered to the load, $V_pI_p = V_sI_s$. When more appliances are switched on across the secondary, the secondary current I_s rises to supply them; and since the primary voltage is fixed by the mains, the primary current I_p must rise in step to keep the input power equal to the larger output power. The transformer therefore draws more current from the mains exactly when more is demanded of its secondary.

13. Eddy currents are the currents induced in any solid piece of metal that sits in a changing magnetic flux. In a transformer the core is bathed in a strongly changing flux, so eddy currents are induced within it and waste energy as heat at the rate $P = I^2R$ in the core itself, which is why they are treated as an unwanted loss. In an induction cooker the same effect is the whole purpose: the changing flux from the coil drives eddy currents in the metal pan, and the pan's own resistance turns those currents into the heat that cooks the food.

EXERCISE 13B

14. Even with no phone connected, the primary winding is still joined across the a.c. mains, so an alternating current flows in it and sets up an alternating magnetic flux in the iron core. This continually changing flux causes hysteresis and eddy-current losses in the core, which appear as heat, while the resistance of the primary winding dissipates a little more. These iron losses continue for as long as the charger is plugged in, whether or not anything is being charged, so the charger keeps drawing a small power from the mains and grows slightly warm.

15. A fluorescent lamp gives out its light in brief bursts, brightest twice in every cycle and almost dark in between, because the power it takes, and so the light it gives, follows the square of the current and therefore peaks one hundred times a second on the 50Hz mains. If a fan blade moves on by one blade-spacing (or a whole number of them) in the time between two flashes, each flash catches a blade in the same place as the one before, so the eye sees the blades frozen; if the timing is slightly out, the blades seem to creep slowly forwards or backwards. The blades are not truly still; they are simply lit only at the instants when they have returned to almost the same position.

16. A running motor is also a generator, for as its coils turn in the magnetic field they generate a back-EMF that opposes the supply voltage, so the current drawn is fixed by the small difference between the supply and the back-EMF. At the moment of switching on the rotor is not yet turning, so there is no back-EMF, and only the low resistance and reactance of the windings limit the current, which is therefore very large. As the motor gathers speed the back-EMF grows and cancels more of the supply voltage, so the current falls to its modest running value.

17. The supply contains a large reservoir capacitor that is charged up to the peak of the rectified voltage while the set is working. A capacitor stores charge, $Q = CV$, and energy, $E = \frac{1}{2}CV^2$, and it can part with them only by driving a current through some conducting path. Once the set is switched off and unplugged there may be no such path, so the capacitor keeps its charge for a long time; if a person then bridges its terminals, the stored charge rushes out through them as a sudden current and delivers a shock.

18. An inductor opposes any change in the current through it by inducing a back-EMF proportional to the rate of change of current, $E = -L\frac{dI}{dt}$. When the switch is opened the current is forced almost instantly to zero, so $\frac{dI}{dt}$ is enormous and a very large back-EMF appears across the widening gap. This high voltage breaks down the air in the gap and drives a bright spark, which is why switching off an inductive load sparks far more fiercely than switching off a plain resistive lamp.

19. The mains completes a fixed number of cycles each second, because the supply authority holds the speed of its generators so steady that the frequency stays at 50Hz on average. A synchronous clock uses a small motor that turns in exact step with these cycles, so that it effectively counts them, and over a day the number of cycles is kept so close to $50 \times 60 \times 60 \times 24$ that the timekeeping is very accurate. Any cycles lost while the frequency dips are made up later when it is allowed to rise, so the long-term count, and hence the time shown, stays correct.

20. A battery can be charged only by a current that flows steadily into its positive terminal, that is by a direct current. If the secondary of the transformer were joined straight to the battery, the alternating current would drive charge into the battery during one half-cycle and draw the same charge out again during the next, so the average current over a cycle would be zero and the battery would gain nothing. A rectifier is therefore needed to turn the alternating current into a one-way current, so that every half-cycle delivers charge into the battery in the same direction.

21. A simple meter of this kind does not measure the rms value directly; it rectifies the waveform, measures its average value, and then multiplies by the form factor of a sine wave, 1.11, so that the scale reads the rms value of a sinusoid. This succeeds only because the mains is very nearly sinusoidal. When the waveform is chopped, as it is by a dimmer, its form factor is no longer 1.11, so the meter's fixed multiplication no longer turns the measured average into the true rms value, and the reading comes out wrong.

22. When a hand grips a live wire carrying the 50Hz mains, the alternating current passes through the muscles and makes them contract; and because the current reverses fifty times a second, the muscles are stimulated again and again and are held in a continuous clench, so the grip tightens on the wire and the person cannot release it. A direct current, by contrast, gives a single steady stimulus that tends to produce one large contraction, which often throws the victim clear of the conductor. This continual gripping is one of the reasons why alternating current at mains frequency is so dangerous.

23. A welding arc carries a very large current at a fairly low voltage, so the welder must turn the 240V mains into a supply of only a few tens of volts but of several hundred amperes. A step-down transformer, with far fewer turns on its secondary than on its primary, lowers the voltage in the ratio of the turns; and since the power is very nearly conserved, $V_p I_p = V_s I_s$, the lower secondary voltage is matched by a correspondingly larger secondary current. It is this large, low-voltage current, carried by the thick secondary wire, that melts the metal to make the weld.

24. The danger from an electric supply comes chiefly from the size of the current it can drive through the body, and for a given body resistance a lower voltage drives a smaller current. In a bathroom or beside a pond, where water lowers the resistance of the skin and makes a shock far more likely, the mains is stepped down by a transformer to a safe 12V, so that any current passing through a person would be small and harmless. The transformer also keeps the low-voltage circuit electrically separate from the mains, which adds a further measure of safety.

25. The appliances in a house are made to run on the 240V, 50Hz alternating supply, and many of them, such as motors, the transformers inside chargers, and fluorescent fittings, will not work on direct current at all. A battery, however, can provide only a low, steady direct voltage, so it cannot drive these appliances by itself. An inverter is therefore used to convert the battery's direct current into an alternating voltage of the right size and frequency, so that the appliances behave just as though they were connected to the mains.

EXERCISE 13C

26. The resistor voltage is in phase with the current, while the inductor and capacitor voltages are a quarter-cycle ahead of and behind it, so the supply voltage is their phasor sum:

$$V = \sqrt{V_R^2 + (V_L - V_C)^2} = \sqrt{(60V)^2 + (80V - 50V)^2} = 67.1V$$

The resistance and the net reactance follow from the common current:

$$R = \frac{V_R}{I} = \frac{60V}{0.50A} = 120\Omega, \qquad X = \frac{V_L - V_C}{I} = \frac{30V}{0.50A} = 60\Omega$$

The supply voltage leads the current by:

$$\phi = \tan^{-1}\frac{V_L - V_C}{V_R} = \tan^{-1}\frac{30V}{60V} = 26.6°$$

and the average power, dissipated only in the resistor, is:

$$P = I^2 R = (0.50A)^2 \times 120\Omega = 30W$$

The supply voltage is 67.1V, the resistance 120Ω, the net reactance 60Ω, the phase angle 26.6°, and the average power 30W.

27. The impedance at each frequency is the voltage divided by the current:

$$Z_1 = \frac{100V}{2.0A} = 50\Omega, \qquad Z_2 = \frac{100V}{1.4A} = 71.4\Omega$$

Since $Z^2 = R^2 + (2\pi fL)^2$ and the second frequency is double the first, the inductive reactance doubles, so subtracting the squared impedances removes R and leaves three times the first reactance squared:

$$3X_1^2 = Z_2^2 - Z_1^2 = (71.4\Omega)^2 - (50\Omega)^2 = 2598\,\Omega^2$$

where $X_1 = 2\pi(50Hz)L$, so:

$$X_1 = \sqrt{\frac{2598\,\Omega^2}{3}} = 29.4\Omega$$

The resistance and inductance then follow:

$$R = \sqrt{Z_1^2 - X_1^2} = \sqrt{(50\Omega)^2 - (29.4\Omega)^2} = 40.4\Omega$$

$$L = \frac{X_1}{2\pi \times 50Hz} = \frac{29.4\Omega}{314\,\text{rad}\,\text{s}^{-1}} = 0.094H$$

The coil has a resistance of about 40Ω and an inductance of about 0.094H.

28. The two sinusoids have different frequencies (50Hz and 150Hz), so the average of their product over a cycle is zero and their mean squares simply add. The rms current is therefore:

$$I_{rms} = \sqrt{\frac{I_1^2 + I_2^2}{2}} = \sqrt{\frac{(3.0A)^2 + (4.0A)^2}{2}} = 3.54A$$

and the average power dissipated in the resistor is:

$$P = I_{rms}^2 R = (3.54A)^2 \times 10\Omega = 125W$$

The combined current has an rms value of 3.54A and dissipates 125W in the resistor.

29. A half-wave rectifier passes only the positive half-cycles. The mean (d.c.) value is that of one hump spread over a whole period:

$$I_{mean} = \frac{I_0}{\pi} = \frac{4.0A}{\pi} = 1.27A$$

Only the half-cycles that are present do any heating, and for a half-wave rectified sinusoid the rms value works out to be half the peak:

$$I_{rms} = \frac{I_0}{2} = \frac{4.0A}{2} = 2.0A$$

so the average power dissipated is:

$$P = I_{rms}^2 R = (2.0A)^2 \times 20\Omega = 80W$$

The mean current is 1.27A, the rms current 2.0A, and the average power 80W.

30. The circuit resonates when the supply frequency equals the natural frequency:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{0.20H \times 2.0 \times 10^{-6}F}} = 252Hz$$

At resonance the reactances cancel and the impedance is just the coil's resistance, so the current is:

$$I = \frac{V}{R} = \frac{10V}{50\Omega} = 0.20A$$

The voltage across the capacitor is the current times its reactance, $X_C = \frac{1}{2\pi f_0 C} = 316\Omega$:

$$V_C = IX_C = 0.20A \times 316\Omega = 63V$$

and the quality factor, the factor by which the capacitor voltage exceeds the supply voltage, is:

$$Q = \frac{1}{R}\sqrt{\frac{L}{C}} = \frac{1}{50\Omega}\sqrt{\frac{0.20H}{2.0 \times 10^{-6}F}} = 6.3$$

The circuit resonates at 252Hz, draws 0.20A, develops 63V across the capacitor, and has a quality factor of 6.3.

31. The quality factor is fixed by the ratio of the resonant frequency to the bandwidth:

$$Q = \frac{f_0}{\Delta f} = \frac{800kHz}{8.0kHz} = 100$$

The capacitance follows from $f_0 = \frac{1}{2\pi\sqrt{LC}}$:

$$C = \frac{1}{(2\pi f_0)^2 L} = \frac{1}{(2\pi \times 8.0 \times 10^5 Hz)^2 \times 2.0 \times 10^{-4}H} = 198pF$$

and the resistance from $Q = \frac{2\pi f_0 L}{R}$:

$$R = \frac{2\pi f_0 L}{Q} = \frac{2\pi \times 8.0 \times 10^5 Hz \times 2.0 \times 10^{-4}H}{100} = 10\Omega$$

Since $\Delta f = \frac{f_0}{Q}$, doubling the resistance halves the quality factor and so doubles the bandwidth to 16kHz. The circuit needs Q = 100, C = 198pF and R = 10Ω.

32. The turns ratio equals the voltage ratio:

$$\frac{N_p}{N_s} = \frac{V_p}{V_s} = \frac{480V}{9600V} = \frac{1}{20}$$

The transmission current is set by the power and the line voltage:

$$I = \frac{P}{V_s} = \frac{24000W}{9600V} = 2.5A$$

so the power wasted as heat in the cable is:

$$P_{loss} = I^2 R = (2.5A)^2 \times 6.0\Omega = 37.5W$$

and the fraction of the generated power reaching the far end is:

$$\frac{24000W - 37.5W}{24000W} \times 100\% = 99.8\%$$

The transformer steps up in the ratio 1:20, the cable carries 2.5A and wastes only 37.5W, so 99.8% of the power is delivered.

33. At full load the output power is the rated apparent power times the power factor:

$$P_{out} = 10000VA \times 0.80 = 8000W$$

Both losses act, so the efficiency is:

$$\eta = \frac{P_{out}}{P_{out} + P_{cu} + P_{Fe}} = \frac{8000W}{8000W + 200W + 80W} = 96.6\%$$

At half load the output halves, but the copper loss falls with the square of the load while the iron loss stays the same:

$$P_{cu}' = 200W \times (0.5)^2 = 50W$$

$$\eta' = \frac{4000W}{4000W + 50W + 80W} = 96.9\%$$

The efficiency is 96.6% at full load and 96.9% at half load, the half-load value being higher because the copper loss has fallen faster than the output.

34. Each branch current is the supply voltage divided by that branch's opposition:

$$I_R = \frac{100V}{50\Omega} = 2.0A, \quad I_L = \frac{100V}{25\Omega} = 4.0A, \quad I_C = \frac{100V}{40\Omega} = 2.5A$$

The resistor current is in phase, the inductor current lags by 90° and the capacitor current leads by 90°, so the line current is:

$$I = \sqrt{I_R^2 + (I_L - I_C)^2} = \sqrt{(2.0A)^2 + (4.0A - 2.5A)^2} = 2.5A$$

lagging the supply voltage by:

$$\phi = \tan^{-1}\frac{I_L - I_C}{I_R} = \tan^{-1}\frac{1.5A}{2.0A} = 36.9°$$

The line current is least when the two reactive branch currents cancel, that is when the capacitor's reactance is made equal to the inductor's, 25Ω; the line current then falls to $I_R = 2.0A$, drawn in phase with the supply. So the branch currents are 2.0A, 4.0A and 2.5A, the line current 2.5A lagging by 36.9°, and the least possible line current is 2.0A.

35. The apparent power is the product of the two meter readings:

$$S = VI = 240V \times 8.0A = 1920VA$$

so the power factor is the ratio of the true power to the apparent power:

$$\cos\phi = \frac{P}{S} = \frac{1200W}{1920VA} = 0.625, \qquad \phi = \cos^{-1}0.625 = 51.3°$$

The reactive power is:

$$Q = \sqrt{S^2 - P^2} = \sqrt{(1920VA)^2 - (1200W)^2} = 1.50 \times 10^3 VAR$$

A purely resistive load of the same true power would draw only $\frac{1200W}{240V} = 5.0A$, so the low power factor forces the supply to carry an extra 3.0A. The power factor is 0.625, the phase angle 51.3°, the reactive power 1.50kVAR, and 3.0A of extra current is drawn.

36. The peak current is $I_0 = 2.0A$, and the greatest energy is stored when the current is at its peak:

$$E_{max} = \frac{1}{2}LI_0^2 = \frac{1}{2} \times 0.50H \times (2.0A)^2 = 1.0J$$

The stored energy depends on i^2, which reaches a maximum twice in every current cycle, so the energy oscillates at twice the supply frequency, that is at $2 \times 50Hz = 100Hz$. The rms current is:

$$I_{rms} = \frac{I_0}{\sqrt{2}} = \frac{2.0A}{\sqrt{2}} = 1.41A$$

and the peak voltage across the inductor, whose reactance is $X_L = 2\pi fL = 157\Omega$, is:

$$V_0 = I_0 X_L = 2.0A \times 157\Omega = 314V$$

The inductor stores at most 1.0J, the energy surging at 100Hz, while the rms current is 1.41A and the peak voltage 314V.

37. At 50Hz the two reactances are:

$$X_L = 2\pi fL = 2\pi \times 50Hz \times 0.30H = 94.2\Omega$$

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50Hz \times 20 \times 10^{-6}F} = 159.2\Omega$$

so the impedance and current are:

$$Z = \sqrt{R^2 + (X_L - X_C)^2} = \sqrt{(100\Omega)^2 + (94.2\Omega - 159.2\Omega)^2} = 119\Omega$$

$$I = \frac{V}{Z} = \frac{240V}{119\Omega} = 2.01A$$

The circuit is net capacitive, so the power factor leads, and the average power is:

$$\cos\phi = \frac{R}{Z} = \frac{100\Omega}{119\Omega} = 0.84, \qquad P = I^2R = (2.01A)^2 \times 100\Omega = 405W$$

The power factor becomes unity at the resonant frequency, where $X_L = X_C$:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{0.30H \times 20 \times 10^{-6}F}} = 65.0Hz$$

where the impedance is R alone, so the current rises to $\frac{240V}{100\Omega} = 2.40A$. At 50Hz the current is 2.01A at a leading power factor of 0.84, dissipating 405W; unity power factor occurs at 65.0Hz, where the current is 2.40A.

38. The turns ratio equals the voltage ratio:

$$\frac{N_p}{N_s} = \frac{240V}{9.0V} = 26.7$$

The peak of the secondary voltage is:

$$V_0 = V_{rms}\sqrt{2} = 9.0V \times \sqrt{2} = 12.7V$$

Full-wave rectification gives a mean (d.c.) output of $\frac{2V_0}{\pi}$:

$$V_{dc} = \frac{2V_0}{\pi} = \frac{2 \times 12.7V}{\pi} = 8.1V$$

and the mean charging current is:

$$I_{dc} = \frac{V_{dc}}{R} = \frac{8.1V}{1.5\Omega} = 5.4A$$

The turns ratio is about 27:1, the secondary peak 12.7V, the d.c. output 8.1V, and the mean charging current 5.4A.

39. The angular speed is $\omega = 2\pi \times \frac{3000}{60} = 314\,\text{rad}\,s^{-1}$ and the area is $A = 200cm^2 = 0.020m^2$, so the peak EMF is:

$$E_0 = NBA\omega = 500 \times 0.040T \times 0.020m^2 \times 314\,\text{rad}\,s^{-1} = 126V$$

The rms EMF and the rms current are:

$$E_{rms} = \frac{E_0}{\sqrt{2}} = \frac{126V}{\sqrt{2}} = 88.9V, \qquad I_{rms} = \frac{E_{rms}}{R} = \frac{88.9V}{100\Omega} = 0.889A$$

and the average power delivered to the resistor is:

$$P = I_{rms}^2 R = (0.889A)^2 \times 100\Omega = 79W$$

The generator gives a peak EMF of 126V (88.9V rms), driving 0.889A and delivering 79W to the resistor.

40. At a power factor of 0.60 the line current is set by the true power, the voltage and the power factor:

$$I_1 = \frac{P}{V\cos\phi} = \frac{30000W}{400V \times 0.60} = 125A$$

so the power wasted in the feeder is:

$$P_1 = I_1^2 R = (125A)^2 \times 0.20\Omega = 3125W$$

After the capacitor bank raises the power factor to 0.90, the line current and feeder loss become:

$$I_2 = \frac{30000W}{400V \times 0.90} = 83.3A, \qquad P_2 = I_2^2 R = (83.3A)^2 \times 0.20\Omega = 1389W$$

so the feeder loss falls by:

$$\frac{3125\text{W} - 1389\text{W}}{3125\text{W}} \times 100\% = 55.6\%$$

Correcting the power factor cuts the line current from 125A to 83.3A and the feeder loss from 3125W to 1389W, a saving of about 56%, all without changing the useful 30kW the workshop consumes.

SOLUTIONS TO EXAMINATION QUESTIONS

PART ONE

Question 1

(a) The error: the atoms in a stretched wire are not moving any faster than at zero strain; the temperature is unchanged. The strain energy is stored as the increase in the interatomic potential energy of the atoms, which sit slightly farther apart than at the equilibrium spacing of the unstretched wire. The work done by the applied force is converted into this potential energy.

(b) Under heavy load each leaf of the spring deflects further, so adjacent leaves slide against each other through a larger distance and the surface friction between them produces the audible squeak. With no load the leaves are in their natural curved shape, the contact between them is essentially static, and no sliding (and therefore no squeak) occurs.

(c)

$$\Delta L = \frac{FL}{AE} = \frac{300\text{N} \times 2.0\text{m}}{1.5 \times 10^{-6}\text{m}^2 \times 2.0 \times 10^{11}\text{Pa}} = 2.0 \times 10^{-3}\text{m}$$

$$U = \frac{1}{2}F\Delta L = \frac{1}{2} \times 300\text{N} \times 2.0 \times 10^{-3}\text{m} = 0.30\text{J}$$

Extension is 2.0mm and strain energy stored is 0.30J.

Question 2

(a) The error: the proportional limit and the elastic limit need not coincide. The proportional limit is the highest stress at which Hooke's law (stress proportional to strain) holds. The elastic limit is the highest stress at which the wire still returns to its original length on unloading. Between these two limits the wire is still elastic but no longer linear in its response. For most metals the two limits are very close but not identical, and the elastic limit lies just above the proportional limit.

(b) Each downstroke loads the rod beyond its elastic limit by a small amount. The rod recovers most of its shape elastically after each stroke but retains a tiny plastic deformation. Over many cycles these tiny plastic deformations accumulate into the visible permanent bend, even though no single stroke applied a force large enough to fracture the rod.

(c)

$$A = \pi\left(\frac{d}{2}\right)^2 = \pi \times (0.5 \times 10^{-3}\text{m})^2 = 7.85 \times 10^{-7}\text{m}^2$$

$$F_{yield} = \sigma_y A = 2.5 \times 10^8\text{Pa} \times 7.85 \times 10^{-7}\text{m}^2 \approx 196\text{N}$$

$$F_{break} = \sigma_b A = 4.0 \times 10^8\text{Pa} \times 7.85 \times 10^{-7}\text{m}^2 \approx 314\text{N}$$

Plastic deformation begins at approximately 196N; fracture occurs at approximately 314N.

Question 3

(a) During an elastic stretch, the applied force builds up gradually from zero to its final value F. The work done is the area under the force-extension graph, which is the area of a triangle: $\frac{1}{2}F\Delta L$. Per unit volume, this becomes $\frac{1}{2}\sigma\varepsilon$ where σ is stress and ε is strain. Multiplying final stress by final strain would assume the force was at its full final value throughout the stretch, which overcounts the work by a factor of two.

(b) The strain energy stored in a Hookean rubber band is proportional to the square of the extension, so doubling the extension stores four times the energy. On release this energy is converted to projectile kinetic energy $\frac{1}{2}mv^2$, so the launch velocity doubles. With doubled launch velocity at the same angle, the projectile range (proportional to v^2) becomes four times its previous value.

(c)

$$U = \frac{1}{2}kx^2 = \frac{1}{2} \times 200\text{Nm}^{-1} \times (0.30\text{m})^2 = 9.0\text{J}$$

$$v = \sqrt{\frac{2U}{m}} = \sqrt{\frac{2 \times 9.0\text{J}}{0.20\text{kg}}} = 9.49\text{ms}^{-1}$$

Energy stored 9.0J; launch velocity approximately 9.5m/s.

Question 4

(a) The isothermal bulk modulus B_T measures the resistance of a fluid to compression when the temperature is held constant; for an ideal gas $B_T = P$. The adiabatic bulk modulus B_S measures the same resistance when no heat flows during the compression; for an ideal gas $B_S = \gamma P$. The two differ because in an adiabatic compression the gas heats up, raising the pressure further than the volume reduction alone produces; the gas appears stiffer when compressed adiabatically.

(b) The bulk modulus, since the speed of sound in a fluid is $v = \sqrt{B/\rho}$. The bulk modulus of water is about $2.2 \times 10^9\text{Pa}$, while that of air at atmospheric pressure is only about $1.4 \times 10^5\text{Pa}$, a ratio of about 10^4. Water molecules are tightly packed and resist compression strongly because the intermolecular separation is already close to the equilibrium value of the strong attractive forces. Air molecules are far apart and offer little resistance to compression beyond the kinetic-pressure term.

(c) (i) Isothermal bulk modulus:

$$B_T = P = 1.5 \times 10^5\text{Pa}$$

(ii) Adiabatic bulk modulus:

$$B_S = \gamma P = \frac{5}{3} \times 1.5 \times 10^5\text{Pa} = 2.5 \times 10^5\text{Pa}$$

Question 5

(a) At every cross-section the same force F is transmitted, but stress is force per unit area; a smaller cross-section therefore carries a larger stress, and the strain $\varepsilon = \sigma/E$ is correspondingly larger. The strain is greatest where the cross-section is narrowest, which is also where the bar is most likely to fail under load.

(b) Glass is a brittle material: under stress it stores elastic energy but cannot dissipate it through plastic deformation. When the elastic limit is reached, the stored energy is released suddenly through the propagation of cracks, shattering the tumbler. A metal cup is ductile: stress beyond the elastic limit causes plastic deformation, in which atoms slip past each other and absorb energy as heat through dislocation motion. The cup dents instead of shattering because plastic flow can absorb the impact energy without catastrophic crack propagation.

(c) Stress in main rod:

$$\sigma_0 = \frac{F}{A_0}$$

Nominal stress at the reduced section:

$$\sigma_{nom} = \frac{F}{A_0/4} = 4\sigma_0$$

Peak stress at the edge of the hole, with stress concentration factor 3:

$$\sigma_{peak} = 3\sigma_{nom} = 12\sigma_0$$

The peak local stress at the hole is 12 times the stress in the unaffected part of the rod.

Question 6

(a) The error: in a parallel-wire cable each wire carries only its share of the total load, not the entire load. Since the wires are identical and forced to extend by the same amount (they are joined at the ends), Hooke's law requires each to carry the same tension. The correct sharing is therefore F/n in each wire.

(b) First, redundancy: if one rope fails, the remaining ropes can still hold the lift (each rope is rated to carry the full load with a safety margin), avoiding a catastrophic fall. Second, fatigue and inspection: thinner ropes flex more easily over the drive sheave, reducing bending stresses and fatigue, and they can be inspected and replaced one at a time without taking the lift out of service for long. A single thick rope would be stiffer, would fatigue faster at the sheave, and would offer no redundancy.

(c) Since the beam is rigid and remains horizontal, all three wires extend by the same amount x. Each tension is $T = (AE/L)x$ for a wire of area A. The middle wire has area 2A, so its tension is twice that of an outer wire.

$$T_1 = T_3 = \frac{AEx}{L}, \quad T_2 = \frac{2AEx}{L}$$

Vertical equilibrium: $T_1 + T_2 + T_3 = W$:

$$\frac{AEx}{L}(1 + 2 + 1) = W \implies \frac{AEx}{L} = \frac{W}{4}$$

$$T_1 = T_3 = \frac{W}{4}, \quad T_2 = \frac{W}{2}$$

Tensions in the ratio $1:2:1$, with the middle (thicker) wire carrying half the load and each outer wire carrying a quarter.

Question 7

(a) Simple harmonic motion. The wire behaves as a Hookean spring of effective force constant $k = AE/L$, where A, L are the wire's cross-sectional area and natural length and E is its Young's modulus. The restoring force on the mass at displacement x from equilibrium is $-kx$, giving SHM with angular frequency $\omega = \sqrt{k/m}$ and period $T = 2\pi\sqrt{m/k}$, which depends on the wire's A, L, and E as well as on the mass.

(b) The suspension acts as a spring-mass system with its own natural frequency, set by the spring constant of the suspension and the sprung mass of the bus. At one particular cruising speed, the rotation frequency of the wheel (or some sub-multiple of it) coincides with this natural frequency, and small unbalanced forces from the wheel drive the suspension into resonance, producing the audible hum. Above and below that speed, the driving frequency is off-resonance and the response is too small to hear.

(c)

$$k = \frac{AE}{L} = \frac{5.0\times10^{-7}\text{m}^2 \times 2.0\times10^{11}\text{Pa}}{1.5\text{m}} = 6.67\times10^{4}\text{Nm}^{-1}$$

$$T = 2\pi\sqrt{\frac{m}{k}} = 2\pi\sqrt{\frac{2.0\text{kg}}{6.67\times10^{4}\text{Nm}^{-1}}} = 3.44\times10^{-2}\text{s}$$

Period of oscillation is approximately 34ms.

Question 8

(a) The proportional limit has been passed: beyond this point the stress is no longer linearly proportional to the strain. Young's modulus is defined for the linear (Hookean) regime, where $\sigma = E\varepsilon$, so the slope of the straight portion of the stress-strain curve gives E directly. The slope of the non-linear portion does not represent any single material constant and cannot be used.

(b) A tensile test loads the sample to failure and records its yield stress, breaking stress and Young's modulus. A below-specification sample fails one or more of these tests directly and is detected. An above-specification sample passes the tension tests, but a bridge cable also fatigues under cyclic loading, suffers stress concentration at fittings, and is exposed to salt corrosion; tensile testing measures only the static-load properties and cannot detect defects that show up only under cyclic, environmental, or geometric stress. Above-specification by tension alone is therefore necessary but not sufficient.

(c)

$$A = \pi\left(\frac{d}{2}\right)^2 = \pi \times (5.0\times10^{-3}\text{m})^2 = 7.85\times10^{-5}\text{m}^2$$

$$\sigma = \frac{F}{A} = \frac{14\times10^{3}\text{N}}{7.85\times10^{-5}\text{m}^2} = 1.78\times10^{8}\text{Pa}$$

$$\varepsilon = \frac{\Delta L}{L} = \frac{0.180\times10^{-3}\text{m}}{0.200\text{m}} = 9.00\times10^{-4}$$

$$E = \frac{\sigma}{\varepsilon} = \frac{1.78\times10^{8}\text{Pa}}{9.00\times10^{-4}} = 2.0\times10^{11}\text{Pa}$$

Young's modulus of the steel is approximately 2.0×10^{11}Pa.

Question 9

(a) The error: a load placed gradually does work $\frac{1}{2}F\Delta L$ on the wire as the force builds up; a load suddenly released does work $F(\Delta L + h)$ as it falls through the wire's stretching distance plus any height h of free fall. The sudden load therefore stores more elastic energy in the wire and causes a larger maximum extension than the equilibrium static extension. After a sudden release the wire oscillates about the static extension, with the first downswing reaching the maximum dynamic extension, which can be much larger than the static value.

(b) In the controlled lowering, the wire stretches by its static extension only and the maximum tension is just the cement weight. In the dropping method, the falling sack converts gravitational potential energy into both elastic strain energy and kinetic energy of the sack just before contact; the wire must then absorb the total energy by stretching beyond its static extension, raising the peak tension well above the static value. The peak tension can exceed the breaking stress of the wire even though the static load alone is well within the safe range.

(c) Effective spring constant of the wire and weight of the load:

$$k = \frac{AE}{L} = \frac{1.0 \times 10^{-6}\text{m}^2 \times 2.0 \times 10^{11}\text{Pa}}{2.0\text{m}} = 1.0 \times 10^{5}\text{Nm}^{-1}$$

$$mg = 5.0\text{kg} \times 9.81\text{ms}^{-2} = 49.05\text{N}$$

Static extension:

$$x_s = \frac{mg}{k} = \frac{49.05\text{N}}{1.0 \times 10^{5}\text{Nm}^{-1}} = 4.91 \times 10^{-4}\text{m} \approx 0.49\text{mm}$$

Energy conservation from release to maximum extension (gravitational PE lost = strain energy stored):

$$mg(h + x_{max}) = \frac{1}{2}kx_{max}^{2}$$

$$\frac{1}{2} \times 1.0 \times 10^{5} x_{max}^{2} - 49.05\, x_{max} - 4.905 = 0$$

Solving the quadratic (taking the positive root):

$$x_{max} = 1.04 \times 10^{-2}\text{m} = 10.4\text{mm}$$

$$\frac{x_{max}}{x_s} = \frac{10.4\text{mm}}{0.49\text{mm}} \approx 21$$

Maximum dynamic extension is approximately 10.4mm, about 21 times the static extension.

Question 10

(a) The fundamental frequency of a stretched wire fixed at both ends is $f = \frac{1}{2L}\sqrt{T/\mu}$, which depends on the tension T, the length L, and the linear mass density μ, but does not contain Young's modulus explicitly. The Young's modulus enters the problem only through the relation between extension and tension; once a particular tension is established in the wire, the wire's vibration frequency is determined by T, L, and μ alone, regardless of which value of E produced that tension. A stiffer wire requires a larger applied force to reach a given tension, but the resulting frequency at that tension is the same as for a softer wire at the same tension.

(b) Tightening the tuning peg increases the tension T in the string. Since $f \propto \sqrt{T}$ at constant L and μ, raising T raises the fundamental frequency, which is heard as a rise in pitch.

(c)

$$f = \frac{1}{2L}\sqrt{\frac{T}{\mu}}$$

$$T = \mu(2Lf)^2 = 8.0 \times 10^{-3}\text{kgm}^{-1} \times (2 \times 0.75\text{m} \times 220\text{Hz})^2 \approx 871\text{N}$$

Tension in the wire is approximately 871N.

Question 11

(a) The property must (i) vary continuously and reproducibly with temperature, so that a single value of the property unambiguously identifies a single temperature; (ii) vary monotonically (single-valued in the working range), so that the same temperature does not correspond to two different values of the property; (iii) vary by an amount large enough to be measured precisely with available instruments, so that small temperature differences produce a distinguishable change in the property.

(b) The reference thermometer is used only at the ice point, the steam point, and the calibration laboratory, so it accumulates very little wear, drift, or chemical contamination. The working thermometer accumulates these effects rapidly through daily handling. By keeping the reference thermometer in a controlled environment, the laboratory preserves a stable standard against which the working thermometer's drift can be detected and corrected.

(c)

$$\theta_{Pt} = \frac{R - R_0}{R_{100} - R_0} \times 100°\text{C}$$

$$\theta_{Pt} = \frac{31.5\Omega - 25.0\Omega}{35.5\Omega - 25.0\Omega} \times 100°\text{C} = 61.9°\text{C}$$

Temperature on the platinum-resistance scale is approximately 62°C.

Question 12

(a) The pressure of an ideal gas at constant volume varies linearly with the absolute temperature, $P/T = \text{const}$, and this relation is independent of the gas used in the limit of low pressure. Different gases (helium, hydrogen, nitrogen) all give the same temperature reading when their thermometers are each calibrated at the triple point of water and used at low enough pressure that they behave ideally. The gas thermometer therefore defines a temperature scale that is essentially independent of the working substance, as required of a thermodynamic temperature scale.

(b) The gas thermometer's calibration is fixed by the laws of physics, not by the specific properties of any one gas, so it does not drift over time and does not need to be re-traced to a higher-tier reference. Platinum resistance thermometers are convenient and quick to read, but their resistance-temperature relation depends on the purity and prior heat-treatment of the platinum, both of which can change with use; they need periodic recalibration against an absolute reference. The gas thermometer fills that role.

(c)

$$T = T_{tp} \times \frac{P}{P_{tp}} = 273.16\text{K} \times \frac{1.099 \times 10^{5}\text{Pa}}{0.800 \times 10^{5}\text{Pa}} = 375.3\text{K} = 102.15°\text{C}$$

Temperature of the bath is approximately 375K, or about 102°C.

Question 13

(a) The error: agreement at two fixed points does not imply agreement at intermediate points. The platinum-resistance scale is constructed by linear interpolation between the ice and steam points, but the actual relationship between the resistance of platinum and its temperature is not exactly linear; the ideal-gas scale, by contrast, is linear in absolute temperature. Wherever the platinum's true response curve deviates from the straight

line that joins the two calibration points, the platinum-resistance scale and the ideal-gas scale disagree, with the largest discrepancy typically at the midpoint of the calibration range.
(b) First, the temperature 400°C is well above the boiling point of mercury (357°C); a mercury-in-glass thermometer cannot operate above this limit without the mercury vaporising and breaking the column. Second, the platinum thermometer can be sealed inside a metal sheath connected to its electronics by long leads, allowing remote reading from outside the steam chamber, while a mercury-in-glass thermometer would have to be visible at the chamber, requiring a window and an operator at the hot location.
(c)

$$\theta_{Pt} = \frac{R - R_0}{R_{100} - R_0} \times 100°C = \frac{41.6 - 25.0}{35.6 - 25.0} \times 100°C = 156.6°C$$

On the ideal-gas scale, $T = 158.0°C$. The platinum-resistance scale reads about 1.4°C below the ideal-gas value at this temperature. The discrepancy is consistent with the slight non-linearity of platinum's resistance-temperature curve over the 0 to 158°C range.

Question 14

(a) At each junction between the two metals, an electric potential difference arises (the Seebeck effect), with magnitude that depends on the temperature of the junction and on the pair of metals. When the two junctions are at different temperatures the two potential differences do not cancel, and a net EMF appears around the loop, driving a current if the loop is closed. The cold junction is necessary because only the temperature difference between the two junctions can be inferred from the EMF; without a reference junction at a known temperature, the hot-junction temperature cannot be deduced from the EMF reading alone.
(b) First, the thermocouple junction is small and has very low thermal mass, so it reaches thermal equilibrium with the meat quickly and does not cool the meat noticeably. Second, the junction can be inserted on a long thin probe into the centre of a piece of meat, while the EMF is read on a meter at a convenient distance; a glass thermometer would have to be inserted along its full length and read at the meat itself.
(c) Linear calibration: $E = k\theta$, with $k = 4.10mV/100°C = 0.0410mV/°C$:

$$\theta = \frac{E}{k} = \frac{7.40mV}{0.0410mV/°C} = 180.5°C$$

Temperature of the oil bath is approximately 180.5°C.

Question 15

(a) The ice point and the steam point are both pressure-dependent: the boiling point of water depends on the surrounding atmospheric pressure, and the melting point of ice changes slightly with pressure. Realising either point to the precision required for primary metrology demands controlling the pressure to high accuracy. The triple point of water, by contrast, is the unique combination of pressure and temperature at which all three phases of water coexist; it is fixed by the substance itself and is independent of external conditions, so a single fixed point is sufficient and is more reproducible than either of the historical pair.
(b) In a triple-point cell, water is sealed inside a glass cell with three phases (ice, liquid water, and water vapour) coexisting at the triple-point temperature 273.16K and pressure 611Pa. So long as all three phases are present, the temperature is fixed by the physics of the substance, independent of the surrounding atmospheric pressure (which cannot reach the sealed water inside the cell). An open ice-water bath, by contrast, is at the ice point at the surrounding atmospheric pressure, which varies from day to day and from one altitude to another; its temperature is therefore less reproducible by a few millikelvins.
(c)

$$T = T_{tp} \times \frac{P}{P_{tp}} = 273.16K \times \frac{28.0kPa}{100kPa}$$

$$T = 273.16K \times 0.280 = 76.5K$$

$$\theta = T - 273.15 \approx -196.7°C$$

Boiling point of liquid nitrogen on the ideal-gas scale is approximately 76.5K or -196.7°C.

Question 16

(a) In a metal, the conduction electrons are essentially free at all temperatures, and the resistance arises mainly from electron scattering by lattice vibrations; rising temperature increases the lattice vibrations and so the resistance grows with temperature. In a semiconductor, the number of conduction electrons is itself temperature-dependent: thermal excitation liberates more charge carriers from the valence band as the temperature rises, and this rapid increase in carrier density dominates the resistance, which therefore falls steeply with temperature. The fall is exponential, so the fractional change in resistance per kelvin is large near room temperature, allowing thermistors to detect small temperature changes that a metal resistance thermometer would resolve only with difficulty.
(b) First, the thermistor's resistance changes more steeply with temperature than that of any practical metal element, so a small temperature change at the patient produces a large change in resistance and hence in the digital display, giving high sensitivity over the narrow clinical range 35 to 42°C. Second, thermistors can be made very small (sub-millimetre), so the probe reaches thermal equilibrium with the patient's tissue rapidly, giving a reading in a few seconds rather than the minute or more needed for a larger-thermal-mass platinum element.
(c) From $R = R_0 e^{(B/T)}$, taking the ratio at the two given temperatures:

$$\frac{R_1}{R_2} = e^{\left(\frac{B}{T_1} - \frac{B}{T_2}\right)} \Rightarrow \ln\frac{R_1}{R_2} = B\left(\frac{1}{T_1} - \frac{1}{T_2}\right)$$

$$B = \frac{\ln R_1/R_2}{1/T_1 - 1/T_2} = \frac{\ln 10.0/3.30}{1/298 - 1/323} = 4270K$$

At $T = 310K$ (37°C):

$$R(37°C) = R_1 e^{\left(B\left(\frac{1}{T} - \frac{1}{T_1}\right)\right)} = 10.0k\Omega \times e^{(-0.554)} = 5.74k\Omega$$

$B = 4270K$ and $R(37°C) = 5.74k\Omega$.

Question 17

(a) A clinical thermometer has a fine constriction in the capillary just above the bulb. As the bulb is heated, mercury expands and is forced through the constriction into the stem, where it rises to the appropriate temperature mark. When the thermometer is removed and cools, the mercury in the bulb contracts and pulls back; but the mercury thread in the stem cannot pass back through the narrow constriction by capillary attraction alone, so it remains in place at the maximum mark. The reading therefore stays at the maximum until the thermometer is shaken to force the mercury back through the constriction.

(b) Shaking accelerates the trapped mercury thread sharply along the stem; the resulting inertial force on the mercury column overcomes the capillary attraction holding it above the constriction, and the column is pushed back through the constriction into the bulb. Without this manual reset, the thermometer would still display the previous patient's maximum and would underread the next patient by however much the next temperature was below the previous reading.

(c)

$$\Delta V = \gamma_{Hg} V_0 \Delta T = 1.82 \times 10^{-4} K^{-1} \times 0.20 cm^3 \times (39.5 - 24.0) K = 5.64 \times 10^{-4} cm^3$$

The volume of mercury increases by approximately $5.6 \times 10^{-4} cm^3$ (about 0.56mm³) over the 15.5K rise. The narrow stem bore converts this small volume change into a visible rise of the mercury thread.

Question 18

(a) Mercury freezes at -39°C and boils at 357°C, giving a working range of about -38 to +350°C; ethanol freezes at -114°C and boils at 78°C, giving a working range of about -110 to +75°C. Mercury therefore covers high temperatures up to several hundred degrees, while alcohol covers very low temperatures (below -38°C) where mercury would solidify. Mercury has a smaller cubic expansion coefficient ($1.82 \times 10^{-4} K^{-1}$) than alcohol ($1.1 \times 10^{-3} K^{-1}$), so for the same bore an alcohol thermometer is roughly six times more sensitive (larger thread movement per kelvin), but covers a narrower range. The choice between mercury and alcohol therefore trades high-temperature reach against low-temperature reach and sensitivity.

(b) -20°C is well above the freezing point of ethanol (-114°C) but already cold enough to be approaching mercury's freezing point (-39°C). On a colder than average night, a mercury thermometer can solidify and stop responding, failing to record the actual minimum. Alcohol remains a free-flowing liquid throughout the temperature range encountered in Mbeya, so an alcohol thermometer is the safe choice for low-temperature work.

(c) Volume change per kelvin of temperature rise:

$$\frac{\Delta V}{\Delta T} = \gamma_{Hg} V_0 = 1.82 \times 10^{-4} K^{-1} \times 0.50 cm^3 = 9.10 \times 10^{-5} cm^3/K$$

This volume change drives a 1.0mm = 0.10cm rise in the stem, occupying a length Δl of bore of cross-sectional area A:

$$\Delta V = A \Delta l \Longrightarrow A = \frac{\Delta V}{\Delta l} = \frac{9.10 \times 10^{-5} cm^3}{0.10 cm} = 9.10 \times 10^{-4} cm^2$$

Cross-sectional area of the stem bore is approximately $9.1 \times 10^{-4} cm^2$ (about 0.091mm², corresponding to a bore diameter of about 0.34mm).

Question 19

(a) Every body at non-zero absolute temperature radiates electromagnetic energy, and for a hot body in the visible-and-infrared range the spectrum and total flux of this radiation depend on the body's temperature, by Planck's law and the Stefan-Boltzmann law. A pyrometer collects this thermal radiation through a lens and infers the temperature of the body from the radiation flux or its spectrum, without any physical contact. Above about 600 to 700°C, mercury-in-glass thermometers fail (mercury boils above 357°C) and platinum resistance thermometers, though usable to about 1000°C, suffer contamination from hot gases and are no longer accurate; for body temperatures of order 1000°C and above, the pyrometer is the only practical instrument.

(b) The molten steel furnace operates at temperatures of order 1500°C, where any contact probe would be destroyed and where the body emits enough visible-and-infrared radiation for a pyrometer to read its temperature easily; the non-contact method is therefore both necessary (for survival of the instrument) and sufficient (for accuracy). The poultry incubator operates at modest temperatures of about 38°C, where a contact probe is survivable, and where a pyrometer would have to be calibrated at low temperatures with much weaker radiation; a thermistor or platinum probe gives a more accurate reading at room-temperature scales.

(c) The pyrometer reading is proportional to T^4:

$$\frac{T^4}{1000^4 K^4} = \frac{4.00}{1.00} \Longrightarrow T = 1000K \times 4^{1/4} \approx 1414K \approx 1141°C$$

Temperature of the glowing body is approximately 1414K, or about 1141°C.

Question 20

(a) Thermal conductivity k is a property of the material alone, with units $Wm^{-1}K^{-1}$. Thermal conductance is a property of a particular component (slab, wall, window), $C = kA/L$ in units WK^{-1}, depending on the material's k, the cross-sectional area A through which heat flows, and the path length L. A wall with a window has two parallel paths for heat: the conductances of the wall and the window add (parallel rule), so the area of each portion enters the total conductance directly. A material of poor insulation occupying a large fractional area can dominate the heat loss, even if its thickness alone is favourable.

(b) The dominant heat-loss path is the window, because the glass window has both a much smaller thickness (5mm vs 100mm of brick) and a small but comparable conductivity to the brick; per square metre, the window conducts much more heat than the brick. Doubling the window area roughly doubles the window contribution, while leaving the small brick contribution unchanged; the overall heat lost therefore increases by approximately the window's share of the total, which can be more than doubling if the original window already carries most of the heat. In practice the increase is close to a doubling because the brick contribution is small to begin with.

(c) (i) Heat conduction through the brick portion:

$$P_{brick} = \frac{k_b A_b \Delta T}{L_b} = \frac{1.0 Wm^{-1}K^{-1} \times 8.5 m^2 \times 10K}{0.10m} = 850W$$

(ii) Heat conduction through the glass window:

$$P_{glass} = \frac{k_g A_g \Delta T}{L_g} = \frac{0.80 Wm^{-1}K^{-1} \times 1.5 m^2 \times 10K}{5.0 \times 10^{-3} m} = 2400W$$

(iii) Total rate of heat conduction:

$$P_{total} = P_{brick} + P_{glass} = 850W + 2400W = 3250W$$

Total heat-loss rate is approximately 3250W, of which the small glass window carries 2400W (about 74%) despite occupying only 15% of the wall area.

Question 21

(a) Heat flowing radially outward through a cylindrical layer crosses successive shells of increasing area $A = 2\pi rL$, so the local temperature gradient dT/dr falls as r increases at the same heat-flow rate P. Integrating Fourier's law $P = -kA(dT/dr)$ between the inner radius r_i and outer radius r_o gives $P = 2\pi kL(T_i - T_o)/\ln(r_o/r_i)$. The relevant geometric factor is therefore the logarithm of the radius ratio, not the thickness $r_o - r_i$ alone; a thin layer wrapped on a large-radius pipe and the same thickness on a small-radius pipe contribute differently to the insulation.

(b) The mineral wool, although a better thermal insulator per unit thickness, is not robust against compression and outdoor weather; the outer foil-faced foam, with a higher conductivity but better mechanical and weather protection, shields the inner layer from damage. The two-layer arrangement therefore combines the high-thermal-resistance of the inner wool with the durability of the outer foam, an option not available with a single-material lagging.

(c) For two cylindrical layers in series carrying the same heat per unit length, the total thermal resistance per unit length is:

$$\frac{R}{L} = \frac{1}{2\pi}\left(\frac{\ln r_2/r_1}{k_1} + \frac{\ln r_3/r_2}{k_2}\right)$$

With $r_1 = 2.0\text{cm}$, $r_2 = 4.0\text{cm}$, $r_3 = 5.0\text{cm}$:

$$\frac{R}{L} = \frac{1}{2\pi}\left(\frac{\ln 2}{0.05} + \frac{\ln 1.25}{0.10}\right) = 2.56\text{KmW}^{-1}$$

$$\frac{P}{L} = \frac{\Delta T}{R/L} = \frac{60\text{K}}{2.56\text{KmW}^{-1}} \approx 23.4\text{Wm}^{-1}$$

Heat loss is approximately 23W per metre of pipe.

Question 22

(a) Newton's cooling law assumes the heat-transfer coefficient h between the body and the fluid is independent of the temperature difference, so that the cooling rate is strictly linear in the excess temperature. In forced convection (a fan or wind), the fluid speed past the body is fixed by external means, h is determined by that speed and is independent of the temperature difference; the law holds rigorously. In natural convection, the buoyancy-driven fluid speed itself depends on the temperature difference (warmer body drives faster updraft), so h increases with ΔT, and the cooling rate depends on ΔT to a power slightly greater than one. At small ΔT the variation in h is small and the law is approximately linear; at large ΔT it is no longer linear.

(b) Fanning forces room-temperature air to flow rapidly past the meat, replacing the warm air layer next to the surface with cool air far more rapidly than buoyancy alone could; the heat-transfer coefficient for forced convection is much larger than for natural convection, so the meat loses heat at a higher rate. Radiation and conduction to the table are essentially unchanged by the fanning, so convection is the dominant mechanism enhanced.

(c) For Newton's law $\theta - \theta_s = (\theta_0 - \theta_s)e^{-kt}$, the cooling constant satisfies $kt = \ln[(\theta_0 - \theta_s)/(\theta - \theta_s)]$.

In free air, $\theta_0 - \theta_s = 60\text{K}$, $\theta - \theta_s = 40\text{K}$, $t = 5.0\text{min}$:

$$k_{\text{free}} = \frac{1}{5.0\text{min}}\ln\frac{60}{40} = 0.0811\text{min}^{-1}$$

In the draught, same temperature change in $t = 1.5\text{min}$:

$$k_{\text{forced}} = \frac{\ln 1.5}{1.5\text{min}} = 0.270\text{min}^{-1}$$

$$\frac{k_{\text{forced}}}{k_{\text{free}}} = \frac{5.0}{1.5} = 3.3$$

Forced-convection cooling constant is approximately 3.3 times the free-convection cooling constant under these conditions.

Question 23

(a) As the temperature of a blackbody rises, the wavelength at which the emission spectrum peaks shifts towards shorter wavelengths according to Wien's displacement law $\lambda_{\text{max}}T = b$, where $b \approx 2.898 \times 10^{-3}\text{mK}$ is the Wien constant. At 1000K the peak lies in the infrared (around 3μm), invisible to the eye; at 2000K it has moved to around 1.5μm, still infrared, but the visible tail of the curve makes the body glow dull red. At 5000K the peak is at 580nm, in the yellow-orange region, and the body glows yellow-white as our Sun does. The colour shift from invisible-infrared through dull red, orange, yellow, and finally bluish-white as temperature rises is a direct visualisation of Wien's law.

(b) The surface temperature of a star sets the wavelength at which its thermal emission peaks, by $\lambda_{\text{max}} = b/T$. Cool red giants have surface temperatures of around 3000K; their peak wavelength is $b/3000 \approx 970\text{nm}$, in the near-infrared, and the visible-spectrum tail is dominated by long red wavelengths. The Sun, at about 5800K, peaks at $b/5800 \approx 500\text{nm}$, in the green-yellow, with a broad visible spectrum that appears whitish-yellow. Hot young stars at 10000K or above peak in the ultraviolet ($b/10000 = 290\text{nm}$), with the visible-spectrum tail weighted to short blue wavelengths, appearing bluish-white. The hierarchy red-yellow-blue corresponds directly to increasing surface temperature.

(c)

$$T = \frac{b}{\lambda_{\text{max}}} = \frac{2.898 \times 10^{-3}\text{mK}}{480 \times 10^{-9}\text{m}} \approx 6038\text{K}$$

For comparison, the Sun's surface temperature from $\lambda_{\text{max}} = 500\text{nm}$:

$$T_{\text{sun}} = \frac{2.898 \times 10^{-3}\text{mK}}{500 \times 10^{-9}\text{m}} \approx 5796\text{K}$$

The star is at approximately 6038K, which is 242K hotter than the Sun's surface.

Question 24

(a) At thermal equilibrium between a body and a cavity at temperature T, the body must emit and absorb radiation at every wavelength at exactly the same rate; otherwise its temperature would not stay equal to T. Since the cavity radiation is a perfect blackbody field, the body's emissivity ε_λ at each wavelength equals its absorptivity α_λ at the same wavelength: this is Kirchhoff's law. A poor emitter at a particular wavelength is therefore necessarily a poor absorber at the same wavelength, and conversely a good emitter at a wavelength is also a good absorber at that wavelength.

(b) During the day, the dominant incoming radiation is solar (peak around 500nm, in the visible), and a matt black surface absorbs strongly at these short wavelengths, heating the roof and the interior more than a reflective white surface would. At night, the roof is itself the warm body radiating to the sky; the dominant outgoing radiation is in the thermal infrared (peak around 10μm at 300K), and the matt black surface emits strongly in the IR by Kirchhoff's law, cooling itself rapidly. The matt black paint therefore **heats** the roof during the day (bad) and **cools** it at night (good); on net, in tropical daytime climates, white or shiny aluminium-coated paint reduces interior temperatures more effectively. The matt black choice is therefore poorly matched to the daytime cooling goal.

(c)

$$\frac{P}{A} = \varepsilon\sigma(T^4 - T_s{}^4)$$

$$\frac{P}{A} = 0.60 \times 5.67 \times 10^{-8}\text{Wm}^{-2}\text{K}^{-4} \times (400^4 - 300^4)\text{K}^4 \approx 595\text{Wm}^{-2}$$

Net radiation loss per unit area is approximately 595Wm^{-2}.

Question 25

(a) The flask consists of a double-walled glass vessel with an evacuated gap between the walls and silvered inner surfaces. The vacuum eliminates both conduction and convection across the gap, since both mechanisms require a material medium. The silvered surfaces have very low emissivity ($\varepsilon \approx 0.05$), reducing radiative heat exchange across the gap by about 20-fold compared with unsilvered glass. Conduction along the narrow neck of the flask cannot be fully eliminated (the lid is in contact with both the contents and the outer wall), and radiation across the silvered gap is reduced but not zero, so all three mechanisms contribute residual heat losses.

(b) Over the eight hours the contents lose only about 20K of temperature, corresponding to an average heat-loss rate of less than a watt for a half-litre of tea (heat capacity around 2.1kJ/K). The slow cooling reflects the success of the vacuum at suppressing conduction and convection across the gap, plus the silvered surfaces at suppressing radiation. The dominant residual loss is conduction along the neck of the flask through the stopper and the threaded fitting, paths that cannot be evacuated; at very long times this neck-conduction loss dominates and limits the flask's effectiveness.

(c) Radiation loss across the silvered vacuum gap (T = 353K, $T_s = 293$K):

$$P_{rad} = \varepsilon\sigma A(T^4 - T_s{}^4) = 0.05 \times 5.67 \times 10^{-8} \times 0.010 \times (353^4 - 293^4) = 0.23\text{W}$$

Free-convection loss from an unlagged cup of the same area:

$$P_{conv} = hA\Delta T = 10\text{Wm}^{-2}\text{K}^{-1} \times 0.010\text{m}^2 \times 60\text{K} = 6.0\text{W}$$

The vacuum-flask radiation loss (0.23W) is approximately 26 times smaller than the open-cup convection loss (6.0W), reflecting the design's effectiveness at suppressing the dominant convective mechanism.

Question 26

(a) A blackbody absorber and a blackbody emitter are necessarily the same surface, by Kirchhoff's law of thermal radiation: $\varepsilon_\lambda = \alpha\lambda$ at every wavelength. A surface that is black at visible-and-near-infrared wavelengths (where the Sun radiates) is automatically also black at the longer thermal-infrared wavelengths (where it radiates at its 350K operating temperature) unless the surface's optical properties differ between the two wavelength ranges. A 'selective surface' achieves this by being coated with a multi-layer film tuned to absorb strongly at solar wavelengths but reflect (and therefore emit poorly) at thermal infrared wavelengths; commercially this is achieved by sputtered metal-on-oxide layers (e.g. nickel-on-tin oxide) and brings emissivity at 10μ m down below 0.1 while keeping absorptivity at 0.5μm above 0.9.

(b) As the panel temperature rises, the rate at which it loses heat by infrared radiation increases as the fourth power of its absolute temperature, while the absorbed solar flux is fixed by the Sun's output. Steady state is reached when the radiative loss plus any convective and conductive losses just equal the absorbed solar power; any further temperature rise would cause net heat loss. The limiting factor on the steady-state temperature is therefore the balance between solar input (constant at noon) and combined infrared and convective output (rising rapidly with T). For a fixed-flux input of 800Wm^{-2} and ordinary panel emissivity, the equilibrium temperature is typically around 70 to 90°C in tropical sunshine.

(c) Solar absorbed:

$$P_{abs} = \alpha SA = 0.90 \times 800\text{Wm}^{-2} \times 2.0\text{m}^2 = 1440\text{W}$$

Radiative loss to surroundings ($T = 343$K, $T_s = 303$K):

$$P_{rad} = \varepsilon\sigma A(T^4 - T_s{}^4) = 0.90 \times 5.67 \times 10^{-8} \times 2.0 \times (343^4 - 303^4) \approx 552\text{W}$$

Net useful collection rate:

$$P_{net} = P_{abs} - P_{rad} = 1440\text{W} - 552\text{W} = 888\text{W}$$

Net useful heat-collection rate is approximately 890W, an efficiency of about 62% relative to the absorbed solar input.

Question 27

(a) On a clear night, the surface radiates strongly in the thermal infrared to a sky whose effective radiative temperature is much lower than the surface temperature (the cold upper atmosphere and ultimately deep space radiate very little back). The radiative loss can exceed the slow conductive heat supply from the warmer air just above and the warmer ground beneath, allowing the surface temperature to fall to the dew point and then below freezing, even when the air at one metre is above 0°C. The local cooling of the surface is fast enough to deposit frost in a few hours of clear sky.

(b) Cloud cover reflects and re-emits thermal radiation back towards the ground, raising the effective sky temperature seen from the surface to a value close to the local air temperature. The net radiative loss $\sigma(T_{ground}{}^4 - T_{sky}{}^4)$ becomes small or even negative, and the surface stays warm enough that frost cannot form. The cloud layer acts as a radiation barrier blocking the surface's view of the cold upper atmosphere; under it, only the much weaker convective cooling operates.

(c)

Clear night, $T_{ground} = 278$K, $T_{sky} = 250$K:

$$\frac{P}{A} = \sigma(T_g^4 - T_{sky}^4) = 5.67 \times 10^{-8} \times (278^4 - 250^4) \approx 117\text{Wm}^{-2}$$

Cloudy night, $T_{ground} = 278$K, $T_{sky} = 280$K:

$$\frac{P}{A} = 5.67 \times 10^{-8} \times (278^4 - 280^4) \approx -10\text{Wm}^{-2}$$

(The negative sign means the ground gains a small amount of net radiation from the warmer cloud layer rather than losing it.)

Under the clear sky the surface loses about 117Wm^{-2} to space, easily enough to drop the surface below 0°C overnight and form frost. Under the cloudy sky the surface gains a small amount of radiation back from the warmer cloud and stays close to the air temperature; no frost forms.

Question 28

(a) The Earth presents a circular cross-section of area πR^2 to the incoming parallel sunlight, so the absorbed power is $(1 - A)S\pi R^2$, where A is the albedo and S the solar constant. The Earth radiates to space from its full spherical surface area $4\pi R^2$, isotropic in direction, with power $4\pi R^2\sigma T^4$ if treated as a blackbody at temperature T. Equating absorbed and emitted powers in steady state gives $\sigma T^4 = (1 - A)S/4$, with the factor 1/4 reflecting the ratio of the cross-sectional intercepting area to the full spherical emitting area.

(b) The Earth's atmosphere is largely transparent to incoming solar radiation (concentrated in the visible-and-near-infrared) but strongly absorbing in the thermal infrared (around 10μm) where the surface radiates. The atmosphere intercepts the surface's outgoing infrared, warms slightly, and re-radiates partly back towards the surface, reducing the net rate of heat loss to space. Steady-state is therefore achieved at a higher surface temperature than for a bare planet without atmosphere; this is the natural greenhouse effect, which raises the surface temperature from the calculated effective value of about -18°C to the observed average of about $+15$°C. Without it, the Earth's surface would be permanently frozen.

(c)

$$\sigma T^4 = \frac{(1-A)S}{4} = \frac{(1-0.30) \times 1370 Wm^{-2}}{4} = 240 Wm^{-2}$$

$$T = \left(\frac{240}{5.67 \times 10^{-8}}\right)^{1/4} \approx 255K$$

Effective equilibrium temperature is approximately 255K, or about $-18°C$, well below the observed surface average of 288K and reflecting the absence of atmospheric greenhouse warming in this calculation.

Question 29

(a) Glass is highly transparent to visible and near-infrared light (roughly 0.4 to 2.5μ m), where solar radiation is concentrated, and admits the bulk of incoming sunlight to the greenhouse interior. At thermal-infrared wavelengths around 10μ m, where a body at about 300K emits its peak radiation, glass is essentially opaque and absorbs strongly. The glass therefore lets sunlight in (warming the interior) but blocks the escape of longer-wavelength thermal radiation from the warmed contents (preventing heat from leaving), a property known as spectral selectivity. The same trick works for a planetary atmosphere with selective greenhouse gases such as CO_2 and H_2O vapour.

(b) Polythene transmits both the incoming visible-and-near-IR radiation **and** the outgoing thermal IR, so it does not block the radiative loss as effectively as glass does. The polythene greenhouse retains heat mainly through its barrier against air-flow convection (it stops the warm interior air from escaping into the outside), while glass adds a radiative-loss barrier on top of the convective barrier. Polythene is therefore a partial improvement over no enclosure (it stops warm-air convection and provides shelter from rain), but its thermal performance is inferior to glass because it does not block the radiative heat-loss pathway.

(c)

$$P_{rad} = \sigma A(T_{glass}^4 - T_{air}^4) = 5.67 \times 10^{-8} \times 4.0 \times (293^4 - 278^4)W \approx 317W$$

A perfectly transparent zero-emissivity film loses zero heat by radiation; the glass panel loses approximately 317W. The radiative loss through actual glass is therefore non-negligible on cool nights, but is the price paid for being opaque to incoming long-wave radiation from a warm interior emitting back to a colder exterior.

Question 30

(a) For an ideal gas, $\frac{1}{2}m\bar{c}^2 = \frac{3}{2}k_B T$, where m is the mass of one molecule and $\bar{c}^2$ the mean square speed. The right-hand side depends only on temperature, so the average translational kinetic energy per molecule is fixed by T alone, independently of the mass of the molecule. Solving for the speed: $v_{rms} = \sqrt{3k_B T/m} = \sqrt{3RT/M}$, which depends inversely on the square root of the molar mass. A lighter molecule moves faster, but with the same average kinetic energy, as a heavier molecule at the same temperature.

(b) Both gases have the same average translational kinetic energy per molecule (set by the common temperature), but the lighter hydrogen molecules ($M = 2gmol^{-1}$) move much faster than the heavier helium molecules ($M = 4gmol^{-1}$). Faster molecules collide with the rubber wall more frequently and have a larger probability of penetrating the microscopic pores of the rubber. Hydrogen therefore effuses out of the balloon faster than helium, and the hydrogen balloon deflates first.

(c)

$$v_{rms,H_2} = \sqrt{\frac{3RT}{M_{H_2}}} = \sqrt{\frac{3 \times 8.314 \times 300}{2.0 \times 10^{-3}}} ms^{-1} \approx 1934 ms^{-1}$$

$$v_{rms,N_2} = \sqrt{\frac{3RT}{M_{N_2}}} = \sqrt{\frac{3 \times 8.314 \times 300}{28 \times 10^{-3}}} ms^{-1} \approx 517 ms^{-1}$$

$$\frac{v_{rms,H_2}}{v_{rms,N_2}} = \sqrt{\frac{M_{N_2}}{M_{H_2}}} = \sqrt{14} \approx 3.74$$

The rms speed of hydrogen at 300K (about $1934 ms^{-1}$) is about 3.74 times that of nitrogen (about $517 ms^{-1}$), reflecting the inverse square-root dependence on molar mass at the same temperature.

Question 31

(a) The total translational kinetic energy of a sample of n moles of an ideal gas is $\frac{3}{2}nRT$, which scales with the number of molecules present. A small flask of gas contains of order 10^{22} molecules; each carries a tiny kinetic energy of order $10^{-21}J$ at room temperature, but the product is of order kilojoules. The gas appears at rest because the molecular velocities are randomly directed, so the vector sum of momenta is zero; the **scalar** kinetic energy is not zero, because energy adds without regard to direction.

(b) The kicked football carries the kinetic energy of a single object of ordinary mass moving at an ordinary speed. The stationary football sits in air, which itself contains an enormous number of molecules per unit volume (of order Avogadro's number per mole, with several moles in even a small volume of surrounding air). Each molecule's translational kinetic energy is tiny, but their sum over the entire surrounding gas vastly exceeds the bulk kinetic energy of any one sports ball. The stationary-football-and-air system therefore contains **more** translational kinetic energy than the kicked football. The gas's energy is distributed isotropically across the molecules, with the vector sum of momenta close to zero; this is why it does not push the football in any particular direction, and the stationary football appears at rest despite being immersed in this large pool of molecular kinetic energy.

(c)

$$KE_{gas} = \frac{3}{2}nRT = \frac{3}{2} \times 2.0mol \times 8.314 Jmol^{-1}K^{-1} \times 300K \approx 7483J$$

$$KE_{ball} = \frac{1}{2}mv^2 = \frac{1}{2} \times 0.45kg \times (25 ms^{-1})^2 \approx 141J$$

$$\frac{KE_{gas}}{KE_{ball}} = \frac{7483J}{141J} \approx 53$$

Total translational KE of the gas is approximately 7483J, about 53 times the kinetic energy of the kicked football. The KE per molecule is tiny, but N_A molecules per mole make the total substantial.

Question 32

(a) **The dashed grey curve**, peaked at higher speed and broadened, corresponds to the higher temperature. As T rises, the average translational kinetic energy per molecule rises in proportion to T, so the distribution is centred on a larger speed: the peak shifts to the right (higher v_p). The total area under each curve must remain equal to the total number of molecules, which is fixed; the peak height therefore falls and the distribution spreads to a wider range of speeds, giving the broader profile at higher temperature.

(b) Although the typical speed of a perfume molecule is around 400ms^{-1} (comparable to a fast bicycle), the molecule does not travel in a straight line through the room: it suffers about 10^{10} collisions per second with surrounding air molecules, and after each collision its direction changes randomly. The net displacement therefore grows much more slowly than the straight-line speed would predict, scaling as $\sqrt{t}$ rather than t (Brownian-like diffusion). The perfume effectively performs a slow random walk, taking many seconds to cross a few metres.

(c) From the diagram, the most probable speed of nitrogen at 300K is approximately 400ms^{-1} (peak position), and the rms speed is approximately 500ms^{-1} (slightly to the right of the peak). From the kinetic-theory formula:

$$v_{rms} = \sqrt{\frac{3RT}{M}} = \sqrt{\frac{3 \times 8.314 \times 300}{28 \times 10^{-3}}}\text{ms}^{-1} \approx 517\text{ms}^{-1}$$

The graph estimate ($\sim 500\text{ms}^{-1}$) is in good agreement with the calculated value (517ms^{-1}).

Question 33

(a) Effusion is the escape of gas molecules through a hole much smaller than the **mean free path** (*the average distance a molecule travels between two successive collisions with other molecules*), so molecules pass one by one without collisions in the hole; the rate is therefore proportional to the mean speed of the molecules and obeys Graham's law strictly. Diffusion is the spreading of one gas through another by random walk, in which molecules undergo many collisions per centimetre travelled; the rate depends on the diffusion coefficient D, which involves both molecular speed and mean free path. Graham's law applies approximately to diffusion (because both processes scale roughly as the molecular speed), but exactly only to effusion.

(b) By Graham's law, the ratio of effusion rates of the two isotopes is $\sqrt{352/349} \approx 1.0043$, a 0.4% enrichment per pass. Starting from natural uranium (0.7% ^{235}U), reaching the 4% enrichment needed for reactor fuel requires the lighter isotope to pass through several hundred porous-membrane stages. Each stage adds only a tiny enrichment, but the cumulative effect of hundreds of stages in series compounds the small per-stage advantage into a useable enrichment.

(c)

$$\frac{r_{He}}{r_{O_2}} = \sqrt{\frac{M_{O_2}}{M_{He}}} = \sqrt{\frac{32}{4.0}} = 2.83$$

Helium effuses at approximately 2.83 times the rate of oxygen under identical conditions, reflecting the inverse-square-root dependence on molar mass.

Question 34

(a) The pitch of a human voice is set by the resonant frequencies of the vocal-tract cavities (mouth, throat), which scale as $f_n \propto v_{sound}/L$, where v_{sound} is the speed of sound in the gas filling the cavity and L is the cavity length. The vocal cords vibrate at the same fundamental frequency in either gas; what changes is the resonant-cavity filtering. Helium has both a higher γ and a much smaller molar mass than air, so v_{sound} in helium is about three times that in air, and all of the formant frequencies of the voice rise by the same factor of about three, producing the audibly higher-pitched effect.

(b) Helium contains no oxygen; inhaling pure helium from a cylinder replaces the oxygen-containing air in the lungs with an inert gas, and the body cannot extract oxygen from helium. Within seconds the oxygen partial pressure in the lung drops to a level too low to support consciousness, and the worker can pass out without warning and (if the pure helium is continued) suffocate. The high-pitched-voice effect is harmless in itself, but the means of producing it (direct inhalation of pure helium) is dangerous because helium does not trigger the body's CO_2-driven breathing reflex; the worker may not feel breathless before passing out.

(c)

$$v_{sound,air} = \sqrt{\frac{\gamma_{air}RT}{M_{air}}} = \sqrt{\frac{1.40 \times 8.314 \times 300}{29 \times 10^{-3}}}\text{ms}^{-1} \approx 347\text{ms}^{-1}$$

$$v_{sound,He} = \sqrt{\frac{\gamma_{He}RT}{M_{He}}} = \sqrt{\frac{(5/3) \times 8.314 \times 300}{4.0 \times 10^{-3}}}\text{ms}^{-1} \approx 1019\text{ms}^{-1}$$

$$\frac{v_{sound,He}}{v_{sound,air}} \approx \frac{1019}{347} \approx 2.94$$

Speed of sound in helium is approximately 1020ms^{-1}, against approximately 347ms^{-1} in air; the helium voice's resonant frequencies are therefore about 2.9 times higher than in air.

Question 35

(a) A diatomic molecule has its two atoms joined by a chemical bond along a single axis. Rotation about an axis perpendicular to this bond carries the two atoms in circles of large radius and stores energy as rotational kinetic energy; rotation about the bond axis itself, however, carries each atom in a circle of essentially zero radius (the atoms sit on the axis), so the moment of inertia about that axis is much smaller than the other two and the rotational quantum is correspondingly larger than k_BT at room temperature. The bond-axis rotational mode is therefore frozen out, and only the two perpendicular rotational modes are accessible at room temperature, giving five active degrees of freedom in total (3 translational + 2 rotational).

(b) By equipartition, each accessible degree of freedom contributes $\frac{1}{2}k_BT$ of energy per molecule, so the total internal energy per molecule is $f \cdot \frac{1}{2}k_BT$. Argon, monatomic with $f = 3$, has total $U = \frac{3}{2}nRT$ per mole. Nitrogen, diatomic with $f = 5$, has total $U = \frac{5}{2}nRT$ per mole. For the same n and T, nitrogen contains 5/3 times as much internal energy as argon, and is correspondingly more capable of doing work or absorbing heat. The same number of molecules can store more energy in the diatomic gas because each molecule has more accessible energy stores.

(c)

$$U_{total} = \frac{f}{2}nRT = \frac{5}{2} \times 1.0 \times 8.314 \times 300J \approx 6236J$$

Translational contribution (f = 3):

$$U_{trans} = \frac{3}{2}nRT = \frac{3}{2} \times 1.0 \times 8.314 \times 300J \approx 3741J$$

Rotational contribution (f = 2):

$$U_{rot} = \frac{2}{2}nRT = nRT = 1.0 \times 8.314 \times 300J \approx 2494J$$

$$\frac{U_{trans}}{U_{total}} = \frac{3}{5} = 60\%, \ \frac{U_{rot}}{U_{total}} = \frac{2}{5} = 40\%$$

Total internal energy is 6236J, of which 60% (3741J) is translational and 40% (2494J) is rotational.

Question 36

(a) At constant volume, the gas does no work on the surroundings, and all heat input goes into increasing the internal energy: $Q_V = \Delta U = (f/2)nR\Delta T$, hence $C_v = (f/2)R$ per mole. At constant pressure, the gas expands as it heats: the heat input raises both the internal energy AND does work $P\Delta V = nR\Delta T$ against the surroundings (using the equation of state). The heat input per mole per kelvin is therefore $C_p = (f/2)R + R = ((f+2)/2)R$, with the additional R corresponding to the work done by the expanding gas at constant pressure. The result $C_p - C_v = R$ is independent of f and is known as Mayer's relation.

(b) Argon is monatomic ($f = 3$) so its $C_v = (3/2)R$, while nitrogen is diatomic ($f = 5$) so its $C_v = (5/2)R$. For the same heating power supplied per mole, the heat-up rate $dT/dt = P_{heat}/(nC_v)$ is larger for argon than for nitrogen, so argon's temperature rises faster but the total energy stored at any temperature is less. In the welding context, this means the argon shielding-gas envelope reaches its design temperature quickly and stays there, with less thermal mass to drag the arc temperature down; the arc actually runs **hotter** in argon, contradicting the premise. The benefit of argon for welding is its inertness (non-reactive with hot metal) rather than thermal-management.

(c)

$$C_v = \frac{f}{2}R = \frac{5}{2} \times 8.314 Jmol^{-1}K^{-1} = 20.79 Jmol^{-1}K^{-1}$$

$$C_p = \frac{f+2}{2}R = \frac{7}{2} \times 8.314 Jmol^{-1}K^{-1} = 29.10 Jmol^{-1}K^{-1}$$

$$\gamma = \frac{C_p}{C_v} = \frac{f+2}{f} = \frac{7}{5} = 1.40$$

For a diatomic ideal gas at room temperature: $C_v = 20.79 Jmol^{-1}K^{-1}$, $C_p = 29.10 Jmol^{-1}K^{-1}$, $\gamma = 1.40$.

Question 37

(a) From the first law $Q = \Delta U + W$, the heat input goes both into raising the internal energy of the gas and into doing work on the surroundings. At constant pressure the gas expands as it heats; the work done by the gas $W = P\Delta V$ is non-zero, so Q exceeds ΔU by exactly W. The relationship is general for any gas-on-piston experiment at constant pressure.

(b) The Sun supplies heat Q to the helium. By the first law $Q = \Delta U + W$, this heat is partitioned between an internal-energy rise ΔU (the helium gets hotter) and the work W done by the gas on the surroundings (the gas expands and pushes back the atmosphere by an amount V at atmospheric pressure P_{atm}). Q exceeds ΔU by $W = P_{atm}\Delta V$, in agreement with the first law of thermodynamics.

(c)

$$\Delta U = nC_v\Delta T = 0.50 \times \frac{3}{2} \times 8.314 \times 300\,J \approx 1871J$$

$$W = nR\Delta T = 0.50 \times 8.314 \times 300\,J \approx 1247J$$

$$Q = nC_p\Delta T = 0.50 \times \frac{5}{2} \times 8.314 \times 300\,J \approx 3118J$$

Heat supplied is approximately 3118J, internal energy change is 1871J, and the work done by the gas is 1247J.

Check: $\Delta U + W = 1871 + 1247 = 3118J = Q$, in agreement with the first law.

Question 38

(a) Internal energy is a state function: it depends only on the current state of the gas (its T, P, V), not on how the gas reached that state. The work done by the gas, by contrast, is $W = \int PdV$ along the chosen path on the PV diagram, which is the area under the chosen path. Different paths between the same two states enclose different areas, so the work done is different. Heat exchange $Q = \Delta U + W$ is also path-dependent, since ΔU is fixed and W varies.

(b) In cool weather the compression is approximately isothermal: heat escapes to the surroundings as the gas is compressed, keeping the temperature close to ambient and the pressure rise more gradual. In hot weather (or when the compression is fast enough that heat cannot escape), the compression is closer to adiabatic: the gas heats during compression and the pressure rises more steeply for the same volume reduction. The work done on the gas equals the area under the compression curve on the PV diagram; the steeper adiabatic curve requires more work to reach the same final pressure than the shallower isothermal curve, so the compressor uses more work in hot weather.

(c) Path 1: isochoric $A \to (P_0/2, V_0)$ has $W_{1a} = 0$. Then isobaric expansion at $P_0/2$ from V_0 to $2V_0$:

$$W_{1b} = \frac{P_0}{2}(2V_0 - V_0) = \frac{1}{2}P_0V_0$$

$$W_{path\ 1} = 0 + \frac{1}{2}P_0V_0 = 0.5\,P_0V_0$$

Path 2: isobaric expansion at P_0 from V_0 to $2V_0$:

$$W_{2a} = P_0(2V_0 - V_0) = P_0V_0$$

Then isochoric drop at $2V_0$: $W_{2b} = 0$:

$$W_{path\ 2} = P_0V_0 + 0 = 1.0\,P_0V_0$$

Path 2 delivers P_0V_0 of work, twice as much as path 1's $0.5\,P_0V_0$, because path 2 expands at the higher pressure P_0 throughout while path 1 expands only at $P_0/2$.

Question 39

(a) The work done by the gas in any infinitesimal step is $dW = PdV$. Around a closed cycle, $W_{net} = \oint P\,dV$, which is the signed area enclosed by the path. The area-as-line-integral rule applies to any closed curve, regardless of shape: rectangular, triangular, oval, or any other. Clockwise traversal on the standard PV diagram gives a positive enclosed area (heat engine direction); counter-clockwise gives a negative area (refrigerator direction). The shape of the cycle determines the area; the direction of traversal determines the sign.

(b) The area enclosed by the indicator diagram on the PV plane equals the net work done by the gas per cycle. The shape of the diagram (triangular, rectangular, or rounded) does not matter for calculating the net work; only the enclosed area does. The technician measures the area (by counting squares on graph paper, by planimeter, or by digital integration) to obtain the work output per cycle, then multiplies by the engine's cycle frequency to get the power output.

(c) The triangle has base $V_C - V_A = 1.0 \times 10^{-3}\text{m}^3$ and height $P_B - P_A = 1.0 \times 10^5\text{Pa}$:

$$W_{net} = \frac{1}{2} \times \text{base} \times \text{height} = \frac{1}{2} \times 1.0 \times 10^{-3}\text{m}^3 \times 1.0 \times 10^5\text{Pa} = 50\text{J}$$

Net work done by the gas per cycle is 50J, equal to the area of the triangular cycle on the PV diagram, with positive sign because the cycle is traversed in the heat-engine direction.

Question 40

(a) For an ideal gas the internal energy depends only on temperature: $U = U(T)$. In an isothermal process, T is constant, so $\Delta U = 0$. The first law $\Delta U = Q - W$ then gives $Q = W$: every joule of heat absorbed from the surroundings is delivered as work done by the gas, with no fraction stored as internal energy. The gas acts as a perfect heat-to-work converter at the surrounding-temperature isotherm, with no temperature change.

(b) The pipeline gas, held by ground contact at the soil temperature, undergoes an essentially isothermal expansion as it flows from high to low pressure. By the first law for an isothermal expansion of an ideal gas, $\Delta U = 0$ and $Q = W$. The gas does positive work $W = nRT\ln(P_i/P_f)$ in expanding against the back pressure; to keep the temperature unchanged, an equal quantity of heat Q must flow into the gas from the surrounding soil. The gas does not visibly heat up because the inflow of heat exactly balances the work loss; the soil cools imperceptibly to compensate.

(c)

$$W = nRT\ln\frac{P_i}{P_f} = 1 \times 8.314 \times 290 \times \ln\frac{80}{5}\,\text{J} \approx 6685\text{J}$$

Work done by the gas during the isothermal expansion is approximately 6685J per mole.

Question 41

(a) Carnot's theorem requires the cycle to be reversible: every infinitesimal step must occur quasi-statically, with no friction, no temperature gradients across heat-exchange surfaces, and no non-equilibrium expansions. Quasi-static means the process is slow enough that the system stays in equilibrium throughout, which implies an infinitely long cycle time and zero power output. Real engines run at finite speed, with finite temperature drops across heat exchangers, friction in moving parts, and finite-rate combustion; all of these introduce irreversibility and reduce the efficiency below Carnot. Practical engines therefore trade some efficiency for finite power output, reaching perhaps 50 to 70% of the Carnot ceiling.

(b) The Carnot ceiling between 150°C (423K) and 27°C (300K) is $\eta_{Carnot} = 1 - 300/423 \approx 29\%$. Two factors drive the realised efficiency below this: (i) finite temperature drops across the heat-exchanger surfaces (the working fluid cannot be at exactly the reservoir temperature; some difference is needed to drive heat flow at the design rate), reducing the effective Carnot ceiling for the working fluid; and (ii) friction and viscous losses in the turbine, pump, and generator, dissipating some of the work output as heat. A realistic geothermal plant operates at 15 to 20% efficiency.

(c)

$$\eta_{Carnot} = 1 - \frac{T_{cold}}{T_{hot}} = 1 - \frac{300\text{K}}{423\text{K}} = 0.291 = 29.1\%$$

$$W_{max} = \eta_{Carnot} \times Q_{hot} = 0.291 \times 1000\text{J} = 291\text{J}$$

Carnot efficiency is approximately 29% for this temperature pair, with at most 291J of work extractable per kilojoule of heat absorbed.

Question 42

(a) The primary mechanism is convection: the warm calorimeter surface heats the air in contact with it, and the warm air rises and is replaced by cooler air. Radiation also contributes, but at the modest excess temperatures of a typical calorimetry experiment (10-30°C above room temperature), convection dominates. The error increases with longer experimental times because the heat loss accumulates: the system loses energy continuously to the surroundings, and the longer it takes for the water and metal to reach thermal equilibrium, the more energy escapes before the final temperature is read.

(b) $COP_{max} = T_{cold}/(T_{warm} - T_{cold}) = 277/(301 - 277) = 277/24 \approx 11.5$. The advertised value 12 slightly exceeds the Carnot ceiling and is therefore physically infeasible. Real refrigerators typically achieve 25 to 50% of the Carnot ceiling, so a realistic COP for these reservoir temperatures is about 3 to 6. The manufacturer's claim of 12 is too good to be true.

(c)

$$COP_{max} = \frac{T_{cold}}{T_{warm} - T_{cold}} = \frac{277\text{K}}{301\text{K} - 277\text{K}} = \frac{277}{24} \approx 11.5$$

$$\frac{Q_{cold}}{dt} = COP_{max} \times P = 11.5 \times 200\text{W} \approx 2308\text{W}$$

Maximum coefficient of performance is approximately 11.5, and the maximum rate of heat extraction is approximately 2300W.

Question 43

(a) After sunset, the iron roof loses heat by two dominant mechanisms. First, radiation: the roof, now at perhaps 50-60°C, radiates infrared energy upward toward the cold night sky (which can be as cold as −20°C in the upper atmosphere). The temperature difference between the roof and the sky is large, driving substantial radiative loss. Second, convection: any breeze passing over the roof carries heat away by forced convection. The roof cools faster than the concrete walls because iron has a much lower thermal mass per unit area than concrete: iron is thin (typically 0.5mm corrugated sheet) and has a relatively low specific heat capacity ($450\text{Jkg}^{-1}\text{K}^{-1}$), while concrete walls are thick (150-200mm) and have a large thermal mass. The roof's small thermal reservoir is emptied quickly by the same rate of radiative and convective loss that would take much longer to cool the massive concrete walls.

(b) The gas does no work on the surroundings as it expands, since the adjacent cylinder is evacuated; there is nothing to push back. The thermal insulation prevents heat exchange with the surroundings. By the first law $\Delta U = 0$, and for an ideal gas this gives $\Delta T = 0$. The temperature of the gas in the combined cylinders is therefore essentially the same as the initial temperature, regardless of the speed at which the expansion

occurred. (For a real gas, intermolecular forces produce a small temperature change in a free expansion: slight cooling for most gases at room temperature, the Joule-Thomson effect, but this effect vanishes for an ideal gas.)
(c) For an ideal gas, free expansion has $Q = 0, W = 0$, hence $\Delta U = 0$ and $\Delta T = 0$:

$$T_f = T_i = 300K$$

Initial volume $V_i = 1.0 \times 10^{-3}m^3$ (half the container); final volume $V_f = 2.0 \times 10^{-3}m^3$. Constant temperature means $PV = const$:

$$P_f = P_i \times \frac{V_i}{V_f} = 1.0 \times 10^5 Pa \times \frac{1}{2} = 5.0 \times 10^4 Pa$$

$$\Delta U = nC_v\Delta T = nC_v \times 0 = 0$$

Final temperature is 300K, final pressure is $5.0 \times 10^4 Pa$ (half the initial value), and the change in internal energy is zero, all in agreement with the qualitative argument from the first law of thermodynamics.

PART TWO

Question 1

(a) At a fixed end the string cannot move. The arriving pulse pulls upward on the support, and by Newton's third law the support pulls downward on the string with an equal force, which launches a reflected pulse on the opposite side, an inversion. At a free end no such reaction force acts: the end is free to overshoot, and the reflected pulse keeps its original orientation. Reflection at a fixed boundary introduces a phase change of π; reflection at a free boundary introduces none.
(b) The projector sends two images to the screen, each carried by light polarised in one of two perpendicular planes. Each lens transmits only the light whose plane of polarisation matches its own transmission axis and blocks the perpendicular component almost completely. So each eye receives only its intended image, and the brain fuses the two slightly different views into one three-dimensional impression.
(c) The wavelength follows from $c = f\lambda$:

$$\lambda = \frac{c}{f} = \frac{3.00 \times 10^8 m/s}{90.0 \times 10^6 Hz} = 3.33m$$

The signal travels at the speed of light, so the time to cover 75km is:

$$t = \frac{d}{c} = \frac{75 \times 10^3 m}{3.00 \times 10^8 m/s} = 2.5 \times 10^{-4} s$$

Question 2

(a) Diffraction, the spreading of a wave as it passes an aperture, becomes appreciable only when the width of the aperture is comparable with the wavelength. The wavelength of light is around $5 \times 10^{-7}m$, vastly smaller than a doorway, so the spreading is utterly negligible and the light appears to travel in straight lines. A slit a fraction of a millimetre wide is far closer to the scale of the wavelength, so the spreading becomes noticeable.
(b) The speed of a water wave in shallow water depends on the depth, falling as the water becomes shallower. As the wave advances up the beach its speed therefore decreases. Its frequency is fixed by the distant source, so since $v = f\lambda$ the wavelength must shorten. The same energy is now packed into a shorter, slower wave, so the crest rises and the wave steepens.
(c) The intensity for a level L is $I = I_0 \times \log^{-1}\left(\frac{L}{10}\right)$. For the two machines:

$$I_1 = (1.0 \times 10^{-12} W/m^2) \times \log^{-1}(8.2) = 1.58 \times 10^{-4} W/m^2$$
$$I_2 = (1.0 \times 10^{-12} W/m^2) \times \log^{-1}(7.8) = 0.63 \times 10^{-4} W/m^2$$

Intensities from independent sources add:

$$I = I_1 + I_2 = 2.21 \times 10^{-4} W/m^2$$

Converting the combined intensity back to a sound level:

$$L = 10\log(2.21 \times 10^8) = 83.5dB$$

Question 3

(a) The range of intensities the ear responds to is enormous: from the faintest audible sound to the threshold of pain spans a factor of about 10^{12}. A linear scale would be hopelessly unwieldy over so vast a range. A logarithmic scale compresses it into a convenient span of roughly 0 to 120. It also matches the ear itself, whose response to intensity is roughly logarithmic, so that equal multiplicative steps in intensity are heard as roughly equal steps in loudness.
(b) The vertical wall reflects the incoming waves with little loss, sending a reflected train back across the path of the incoming train. At every point the water carries both trains at once, and by the principle of superposition the resultant displacement is the sum of the two. Because the trains travel in different directions, their crests and troughs overlap in a continually shifting pattern, giving the choppy criss-cross appearance.
(c) The grating spacing is the reciprocal of the line density:

$$d = \frac{1mm}{400} = 2.5 \times 10^{-6} m$$

The grating equation is $d\sin\theta = m\lambda$. The largest order corresponds to $\sin\theta = 1$:

$$m_{max} = \frac{d}{\lambda} = \frac{2.5 \times 10^{-6} m}{600 \times 10^{-9} m} = 4.17$$

The order must be a whole number, so the highest maximum that can be observed is the **fourth** order.

Question 4

(a) While the pulses overlap, the displacement of each point of the string is the algebraic sum of the displacements the two pulses would separately produce, which is the principle of superposition. The medium responds to each pulse independently, so neither is altered by the other. Once they have passed through each other, each pulse continues exactly as before, since no energy has been exchanged between them.
(b) The gun emits microwaves of known frequency, which reflect from the moving car and return. Because the car is moving, the reflected waves are Doppler-shifted: the car first acts as a moving receiver and then as a moving re-emitter, so the echo is shifted twice. The gun measures the frequency difference between the emitted and returned waves; this difference is proportional to the car's speed, from which the speed is found.
(c) By Malus's law the transmitted intensity is $I = I_0(\cos\theta)^2$. Setting $I = \frac{1}{4}I_0$:

$$(\cos\theta)^2 = \frac{1}{4}$$

$$\cos\theta = \frac{1}{2}$$

The required angle is therefore:

$$\theta = 60°$$

Question 5

(a) Sound is a mechanical wave: it travels only by the successive disturbance of the particles of a material medium, each layer pushing on the next. The Moon has no atmosphere, so between the shattering rock and the astronaut there is no medium to carry the compressions and rarefactions. With nothing to vibrate, no sound wave can form, and the event reaches the astronaut as silence.

(b) The colours come from interference, not pigment. The surfaces carry extremely fine, regularly spaced structures whose spacing is comparable with the wavelength of light. Light reflected from successive layers interferes; for a given viewing angle one wavelength is reinforced while others are cancelled, so a colour is seen. As the viewing angle changes, the path difference changes, the reinforced wavelength changes, and the colour shifts. This is structural colour.

(c) The wave speed is set by the tension and the linear density:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{80N}{0.020kg/m}} = 63.2m/s$$

The average power carried by the wave is $P = \frac{1}{2}\mu\omega^2 A^2 v$, with $\omega = 2\pi f = 314rad/s$. Substituting:

$$P = \frac{1}{2}(0.020kg/m)(314rad/s)^2(0.0080m)^2(63.2m/s) = 3.99W$$

Question 6

(a) The speed of a transverse wave on a string is $v = \sqrt{T/\mu}$: it is fixed entirely by the tension and the mass per unit length, properties of the string itself, and is independent of how the wave is made. The frequency is the rate at which the source shakes the end of the string, each oscillation launching one cycle, so the source alone fixes it. The wavelength then adjusts to satisfy $\lambda = v/f$.

(b) The regularly spaced openings of the mesh act as a two-dimensional diffraction grating. Light from the lamp passing through the many openings is diffracted, and the diffracted beams interfere. In directions where the path difference between neighbouring openings is a whole number of wavelengths the interference is constructive, giving a bright spot; the regular spacing of the mesh produces a regular pattern of such spots.

(c) A pipe closed at one end has a node at the closed end and an antinode at the open end, so the fundamental wavelength is four times the length:

$$\lambda_1 = 4L = 4(0.55m) = 2.20m$$
$$f_1 = \frac{v}{\lambda_1} = \frac{340m/s}{2.20m} = 154.5Hz$$

A closed pipe resonates only in the odd harmonics, so the next two resonances are the third and fifth harmonics:

$$f_3 = 3f_1 = 463.6Hz$$
$$f_5 = 5f_1 = 772.7Hz$$

Question 7

(a) In a double slit only two beams interfere, so the intensity varies gradually with angle and the bright fringes are broad. In a grating, light from many thousands of slits interferes. At an exact maximum every beam arrives in phase. A very slight turn away from that direction is enough for the many beams, spread across the whole grating, to fall progressively out of step and cancel completely. The maxima are therefore confined to very narrow ranges of angle, and the more slits there are, the sharper the line.

(b) The echo returning from the moth is Doppler-shifted. If the moth is flying towards the bat it acts as an approaching reflector and the echo returns at a higher frequency than the emitted pulse; if it is flying away, the echo returns at a lower frequency. By comparing the pitch of the echo with that of its own pulse, the bat senses whether the moth is closing in or escaping.

(c) For two waves of the same frequency the resultant amplitude is:

$$A_R = \sqrt{{A_1}^2 + {A_2}^2 + 2A_1A_2\cos\varphi}$$

Substituting $A_1 = 3.0cm$, $A_2 = 5.0cm$ and $\varphi = 60°$:

$$A_R = \sqrt{(3.0cm)^2 + (5.0cm)^2 + 2(3.0cm)(5.0cm)(0.50)}$$
$$A_R = 7.0cm$$

Question 8

(a) In a stationary wave every particle oscillates with the same frequency, but the amplitude varies with position, being zero at the nodes and greatest at the antinodes. Within one loop all the particles pass through equilibrium together and reach their extremes together: they oscillate exactly in phase, differing only in amplitude. Across a node the sign of the displacement reverses, so particles in adjacent loops are always half a cycle out of phase, reaching their extremes on opposite sides at the same instant.

(b) When the jet moves faster than the speed of sound it outruns the waves it produces. These waves cannot move ahead of it; instead they pile up and overlap along a cone-shaped shock front trailing the aircraft, across which there is a sharp, intense pressure change. When this shock front sweeps over an observer on the ground, the abrupt pressure change is heard as the single loud bang.

(c) The first dark fringe of a single slit occurs where $a\sin\theta = \lambda$:

$$\sin\theta = \frac{\lambda}{a} = \frac{600\times10^{-9}m}{0.20\times10^{-3}m} = 3.0\times10^{-3}$$

So $\theta = 3.0\times10^{-3}rad$.

The central band runs from this fringe on one side to the matching fringe on the other, so on a screen at distance D its width is:

$$w = 2D\sin\theta = 2(2.0m)(3.0\times10^{-3}) = 1.2\times10^{-2}m$$

The central bright band is therefore 12mm wide.

Question 9

(a) In two-slit interference the two beams have effectively equal amplitude wherever they meet, so wherever they arrive in phase they reinforce to the same maximum brightness, and all the bright fringes are alike. In single-slit diffraction the slit acts as very many tiny sources across its width. At the centre all their contributions arrive in phase and add fully; away from the centre they are increasingly out of step, so only a partial reinforcement survives, and the bright bands fade with distance from the centre.

(b) Speech sets the base of the first tin vibrating; the vibration passes into the string and travels along it as a wave to the second tin, whose base then re-radiates the sound. A string can carry such a wave only while it is taut, because the wave needs tension to supply the restoring force, its

speed being $v = \sqrt{T/\mu}$. When the string goes slack the tension falls almost to zero, the string can no longer transmit the vibration, and the sound fails to cross.

(c) The sound is Doppler-shifted twice: once as the moving car sends sound to the cliff, and again as the moving driver receives the reflected sound. For a source and observer moving towards a stationary reflector at the same speed u, the two shifts combine to give:

$$f' = f\frac{v+u}{v-u}$$

$$f' = 480\text{Hz} \times \frac{340\text{m/s} + 30\text{m/s}}{340\text{m/s} - 30\text{m/s}} = 573\text{Hz}$$

The driver therefore hears the echo at:

$$f' = 573\text{Hz}$$

Question 10

(a) The flash and the thunder are produced together, but they reach the observer at very different speeds. Light travels at about 3×10^8m/s, so the flash arrives almost instantly. Sound travels through air at only about 340m/s, nearly a million times slower, so the thunder takes an appreciable time to cover the same distance. The further away the stroke, the longer the gap between the flash and the thunder.

(b) Heating in air grows a very thin, transparent film of oxide on the steel, and the film thickens the longer and hotter the heating. Light reflected from the top of the oxide film and light reflected from the steel beneath it interfere. The film thickness fixes the path difference, and so fixes which wavelength is reinforced and which cancelled, so the surface shows an interference colour. As the film thickens the reinforced wavelength changes, carrying the surface through the observed sequence of colours.

(c) The fundamental frequency of a wire fixed at both ends is $f = \frac{1}{2L}\sqrt{T/\mu}$. With the length and the linear density unchanged, $f \propto \sqrt{T}$, and therefore $T \propto f^2$.

The frequency is to rise in the ratio $f_2/f_1 = 3/2$, so:

$$\frac{T_2}{T_1} = \left(\frac{f_2}{f_1}\right)^2 = \left(\frac{3}{2}\right)^2 = 2.25$$

The tension must be increased to 2.25 times its original value.

Question 11

(a) Both systems, when displaced from equilibrium, experience a restoring force directed back towards the equilibrium position and proportional to the displacement from it. For the spring it is the elastic force $F = -kx$; for the pendulum at small angles it is the tangential component of gravity, which is likewise proportional to the displacement. Whenever the restoring force is proportional to the displacement and oppositely directed, the acceleration obeys $a = -\omega^2 x$, which is the defining condition of simple harmonic motion. The two systems differ in their physical make-up but share this one essential property.

(b) Where the protector and the screen nearly touch there is a very thin film of trapped air whose thickness varies from point to point. Light reflected from the lower surface of the protector and light reflected from the screen beneath it interfere. Where the air gap makes the path difference a whole number of wavelengths the interference is constructive for that colour, and where it makes an odd number of half-wavelengths it is destructive; since the gap thickness varies across the contact region, different colours are reinforced at different places, producing the coloured rings and bands. This is the same effect as Newton's rings.

(c) The observer moves towards a stationary source, so the observed frequency is raised. With the observer's speed u_o added to the speed of sound:

$$f' = f\frac{v+u_o}{v}$$

$$f' = 660\text{Hz} \times \frac{(340\ +\ 20)\text{m/s}}{340\text{m/s}} = 660\text{Hz} \times \frac{360\text{m/s}}{340\text{m/s}}$$

The motorcyclist therefore hears:

$$f' = 699\text{Hz}$$

Question 12

(a) The fringes form a thin-film interference pattern, and a given fringe is the locus of all points at which the air gap between the curved lens and the flat plate has one particular thickness. The lens surface is part of a sphere resting on a plane, so the air gap has the same thickness at every point lying at a given distance from the point of contact. Points of equal gap thickness therefore lie on circles centred on the contact point, and so each fringe, bright or dark, is itself a circle.

(b) Springing off leaves the board both displaced and moving, so it has been given elastic potential energy and kinetic energy. With the diver gone, no further force drives it; the board simply oscillates freely at its own natural frequency, energy passing back and forth between the elastic strain of the bent board and the kinetic energy of its motion. This is a free oscillation, and it dies away only as the energy is gradually lost to air resistance and internal friction.

(c) For a point source the intensity falls with the inverse square of the distance:

$$\frac{I_2}{I_1} = \left(\frac{r_1}{r_2}\right)^2 = \left(\frac{2.0\text{m}}{8.0\text{m}}\right)^2 = \frac{1}{16}$$

The corresponding change in sound level is:

$$\Delta L = 10\log\left(\frac{I_2}{I_1}\right) = 10\log\left(\frac{1}{16}\right) = -12.0\text{dB}$$

The level at 8.0m is therefore:

$$L_2 = 88\text{dB} - 12.0\text{dB} = 76.0\text{dB}$$

Question 13

(a) An open pipe has a displacement antinode at each end, so its fundamental fits half a wavelength into the pipe: the fundamental wavelength is twice the pipe length. A pipe closed at one end has a node at the closed end and an antinode at the open end, so its fundamental fits only a quarter of a wavelength into the same length: its fundamental wavelength is four times the pipe length, twice that of the open pipe. Since frequency is inversely proportional to wavelength, the closed pipe sounds at half the frequency, which is one octave lower.

(b) The baby and the elastic strap together form a mass-on-a-spring system. The strap stretches under the baby's weight to an equilibrium position; if the baby is displaced below it the extra tension provides a restoring force proportional to the displacement, and above it the reduced tension does

the same. The restoring force being proportional to displacement, the baby executes simple harmonic motion, with a regular period $T = 2\pi\sqrt{m/k}$ set by the baby's mass and the stiffness of the strap, hence the steady rhythm.

(c) The fringe spacing is $y = \frac{\lambda D}{d}$, so it is proportional to the wavelength of the light. On immersion in water the wavelength shortens to $\lambda_w = \lambda/n$, the frequency being unchanged, so the fringe spacing shrinks in the same ratio:

$$y_w = \frac{y}{n} = \frac{1.8\text{mm}}{1.33} = 1.35\text{mm}$$

Question 14

(a) In simple harmonic motion the maximum speed is $v_{max} = \omega A$ and the maximum acceleration is $a_{max} = \omega^2 A$; both are directly proportional to the amplitude, so doubling the amplitude doubles each of them. The period $T = 2\pi/\omega$ is fixed by the angular frequency alone, and ω is set by the properties of the system, the stiffness and mass or the length and gravity, not by how far the oscillator is displaced. Doubling the amplitude does not change ω, so the period is unchanged: the oscillator simply covers twice the distance each cycle, moving twice as fast to do so in the same time.

(b) The smooth curved wall of the dome reflects sound well. Sound spreading out from the whisper at one point strikes the curved surface and is reflected, and the curvature is such that the reflected sound is brought together again, or focused, at the second point. The whisper's energy, instead of spreading thinly in every direction, is concentrated by these reflections onto that one listening point, so it arrives there loud enough to hear, while points away from the focus receive only the faint direct sound.

(c) Light reflected from the top surface of the film undergoes a phase change of π, while light reflected from the bottom surface does not. Constructive interference in the reflected light, giving a strongly reflected colour, therefore requires:

$$2nt = \left(m + \frac{1}{2}\right)\lambda$$

The optical path is $2nt = 2(1.33)(1.10 \times 10^{-7}\text{m}) = 2.93 \times 10^{-7}\text{m}$. Taking the lowest order, $m = 0$:

$$\lambda = 2 \times (2.93 \times 10^{-7}\text{m}) = 5.85 \times 10^{-7}\text{m}$$

The most strongly reflected visible wavelength is 585nm.

(The next order, m = 1, would give 195nm, which lies in the ultraviolet.)

Question 15

(a) The speed of sound in a gas depends on how rapidly the molecules pass the disturbance from one to the next. In warmer air the molecules move with greater average speed, so they collide more frequently and transmit the compressions and rarefactions more quickly. The speed of sound in a gas is in fact proportional to the square root of the absolute temperature, so a rise in temperature raises the speed.

(b) Each fine scratch is a narrow transparent track of a width not enormously greater than the wavelength of light, and it diffracts the light passing through it, spreading that light out in the direction across the scratch. The bright point of a distant headlight is therefore smeared into a streak of light running perpendicular to the scratch. Because the wiper scratches lie along many arcs at different orientations, the diffracted light spreads in a range of directions, producing the spikes and streaks seen around each headlight.

(c) For a mass on a spring $T = 2\pi\sqrt{m/k}$, so the period is inversely proportional to the square root of the stiffness. Two identical springs side by side, sharing the load, act as a single spring of twice the stiffness:

$$k_p = 2k$$

The period therefore shortens by a factor of $\sqrt{2}$:

$$T_2 = \frac{T_1}{\sqrt{2}} = \frac{0.80\text{s}}{\sqrt{2}} = 0.57\text{s}$$

Question 16

(a) At the extremes of the motion the oscillator is momentarily at rest, so all its energy is potential. For a system with restoring force $F = -kx$, the potential energy stored at displacement x is $\frac{1}{2}kx^2$, and at the extreme, where $x = A$, this is $\frac{1}{2}kA^2$. Since the total energy is constant throughout the motion, in the absence of damping, and equals this maximum potential energy, the total energy is $\frac{1}{2}kA^2$, proportional to the square of the amplitude. Doubling the amplitude therefore quadruples the energy of the oscillation.

(b) The wire is a narrow obstacle of width comparable with the wavelength of light. As light from the distant lamp passes the wire it is diffracted, bending into the region behind the wire. Light arriving at a point in that region by the two paths passing the opposite sides of the wire has a path difference that depends on the direction; where the path difference is a whole number of wavelengths the waves reinforce, and where it is an odd number of half-wavelengths they cancel. This alternation produces the regular pattern of light and dark bands.

(c) An open pipe has a displacement antinode at each end, so its fundamental wavelength is twice the length of the pipe, $\lambda_1 = 2L$. Hence:

$$L = \frac{v}{2f_1} = \frac{340\text{m/s}}{2 \times 256\text{Hz}} = 0.664\text{m}$$

An open pipe sounds all the harmonics, so the second harmonic has twice the fundamental frequency:

$$f_2 = 2f_1 = 2 \times 256\text{Hz} = 512\text{Hz}$$

Question 17

(a) For a steady interference pattern the light reaching the two slits must be coherent: the two slits must always be illuminated by waves with a constant phase relationship. An ordinary lamp is an extended source whose different parts emit independently, so light reaching the two slits from it directly would have no fixed phase relationship. The single narrow slit acts as one small, effectively point-like source: the light spreading from it reaches both slits as parts of the same wavefront, so the two slits are illuminated coherently and a stable pattern is formed.

(b) When a person speaks, the sound reaches their own ears by two routes, through the air and directly through the bones and tissues of the head. Bone conduction carries the lower frequencies particularly well, so the voice a speaker normally hears is enriched in low frequencies and sounds fuller and deeper to them. A recording captures only the air-conducted sound; played back, it reaches the ears by the air route alone, without the bone-conducted bass, so it sounds thinner and unfamiliar to the speaker, though it is the voice that others always hear.

(c) The period depends only on the mass and the stiffness:

$$T = 2\pi\sqrt{\frac{m}{k}} = 2\pi\sqrt{\frac{0.30\text{kg}}{30\text{N/m}}} = 0.63\text{s}$$

At the equilibrium position the spring tension balances the weight, $mg = 0.30kg \times 9.8m/s^2 = 2.94N$. During the motion the tension varies above and below this by $kA = 30N/m \times 0.080m = 2.4N$. The greatest tension, at the lowest point, is:

$$F_{max} = mg + kA = 2.94N + 2.4N = 5.34N$$

And the least tension, at the highest point, is:

$$F_{min} = mg - kA = 2.94N - 2.4N = 0.54N$$

Question 18

(a) The displacement completes one full cycle in each period. The kinetic energy depends on the square of the speed. In one period of the displacement the oscillator passes through the equilibrium position twice, once moving each way, and at each passage the speed, and hence the kinetic energy, reaches its maximum; the kinetic energy falls to zero twice as well, at the two extremes. The kinetic energy therefore rises and falls twice in every single cycle of the displacement, so it varies at twice the frequency.

(b) The chest-piece picks up the heartbeat sounds, and the hollow tubes of the stethoscope carry the sound waves directly to the doctor's ears. Confined within the narrow tubes, the sound energy cannot spread out and weaken as it would in the open air, where it would diminish with the square of the distance. Almost all of the collected sound energy is therefore delivered to the ears instead of being dispersed, so the heartbeat is heard far more loudly than it would be by simply listening near the chest.

(c) The grating spacing is:

$$d = \frac{1mm}{500} = 2.0 \times 10^{-6}m$$

From the grating equation $d\sin\theta = m\lambda$, with $m = 2$ and $\theta = 36.0°$:

$$\lambda = \frac{d\sin\theta}{m} = \frac{(2.0 \times 10^{-6}m)(\sin 36.0°)}{2}$$

$$\lambda = 5.88 \times 10^{-7}m$$

The wavelength of the light is 588nm.

Question 19

(a) When light strikes a surface, the component of its vibration parallel to the surface is reflected more strongly than the component lying in the plane of incidence. The reflected beam is therefore richer in the parallel vibration, that is, partially polarised. At one particular angle of incidence, Brewster's angle, the reflected and refracted rays are at right angles to each other; the vibration in the plane of incidence is then not reflected at all, and the reflected light consists solely of the vibration parallel to the surface, so it is completely polarised.

(b) The hanging load behaves as a simple pendulum, with the cable serving as the pendulum's length. The period of a simple pendulum is $T = 2\pi\sqrt{l/g}$: it increases as the square root of the length. Paying out more cable lengthens the pendulum, so the period grows and each swing takes longer, the load swinging more slowly. The mass of the load does not enter the formula, so it is the cable length alone that the operator alters when the swinging becomes lazier.

(c) The pulse travels out to the aircraft and back, so in the measured time it covers twice the distance to the aircraft:

$$2d = ct$$

$$d = \frac{ct}{2} = \frac{(3.00 \times 10^8 m/s)(1.20 \times 10^{-4}s)}{2}$$

The aircraft is therefore:

$$d = 1.80 \times 10^4 m = 18.0km$$

Question 20

(a) All three are electromagnetic waves: each consists of oscillating electric and magnetic fields, at right angles to each other and to the direction of travel, and each is produced by accelerating electric charges. They all travel through a vacuum at the same speed, 3×10^8m/s, and all show the wave behaviours of reflection, refraction, interference, diffraction and polarisation. They differ only in frequency, and so in wavelength; this difference gives them their different uses and effects, but their common nature makes them one single family, the electromagnetic spectrum.

(b) Toughened glass is made by cooling its surfaces rapidly, which leaves the glass with permanent built-in stresses that vary from place to place. Stressed glass is birefringent: it affects the two perpendicular components of a light wave differently, altering the polarisation of light passing through it by an amount that depends on the local stress. The daylight reaching the window is already partly polarised, and the sunglasses act as an analysing polariser. Where the stress in the glass has turned the polarisation away from the sunglasses' axis, less light passes and a dark blotch is seen; the varying stress thus shows itself as a pattern of light and dark patches.

(c) From the period formula $T = 2\pi\sqrt{l/g}$, the length is:

$$l = \frac{gT^2}{4\pi^2} = \frac{(9.8m/s^2)(2.0s)^2}{4\pi^2} = 0.993m$$

Since $T \propto \sqrt{l}$, a small fractional change in length produces half that fractional change in period:

$$\frac{\Delta T}{T} = \frac{1}{2}\frac{\Delta l}{l} = \frac{1}{2}(1.0\%) = 0.5\%$$

The period increases by 0.5%.

Question 21

(a) The natural frequency of a system is the frequency at which it oscillates on its own once it has been displaced and released, with no external driving force acting. It is fixed entirely by the properties of the system. For a mass on a spring it is set by the force constant of the spring and the mass, through $f = \frac{1}{2\pi}\sqrt{k/m}$. For a simple pendulum it is set by the length of the pendulum and the acceleration of free fall, through $f = \frac{1}{2\pi}\sqrt{g/l}$.

(b) Each spectator only rises and sits down again about their own seat; the disturbance, the act of standing, is passed on from each spectator to the next, so the pattern of standing sweeps around the ground. What travels round the stadium is the disturbance and the energy carried with it, while the spectators, who form the medium, merely oscillate in place. This is the defining feature of a wave: the wave profile and its energy move forward, the medium itself does not.

(c) The clap travels from the girl to the cliff and back, covering twice her distance from the cliff, so:

$$d = \frac{vt}{2} = \frac{(340m/s)(1.5s)}{2} = 255m$$

After walking 60m closer, her new distance from the cliff is $255m - 60m = 195m$. The echo now covers twice this distance:

$$t' = \frac{2d'}{v} = \frac{2(195m)}{340m/s} = 1.15s$$

Question 22

(a) Sound is carried forward by each particle of the medium disturbing its neighbours, and the speed depends on how stiffly the particles are coupled and on the density of the medium. In a solid or a liquid the particles are much more closely bound, and the elastic restoring forces between them far stronger, than in a gas, so a disturbance is handed on from particle to particle far more rapidly. Solids and liquids are also denser, which on its own would slow the sound, but the great increase in stiffness outweighs the increase in density, so the net result is a much higher speed.

(b) On the stretched cord the jumper oscillates up and down: below the equilibrium extension the cord pulls her up, above it her weight pulls her down, giving a restoring force and so an oscillation that is close to simple harmonic. The motion is damped. At each bounce some mechanical energy is lost, to air resistance and to internal friction within the stretching cord, and is carried away as heat. With less energy in each successive cycle the amplitude steadily falls, bounce by bounce, until the surplus energy is spent and she hangs at rest at the equilibrium extension.

(c) The wavelength follows from $v = f\lambda$:

$$\lambda = \frac{v}{f} = \frac{1.8m/s}{2.5Hz} = 0.72m$$

One whole wavelength corresponds to a phase difference of 2π radians, so two points a distance d apart differ in phase by:

$$\Delta\varphi = 2\pi\left(\frac{d}{\lambda}\right) = 2\pi\left(\frac{0.30m}{0.72m}\right)$$

which gives:

$$\Delta\varphi = \frac{5\pi}{6}rad = 2.6rad$$

Question 23

(a) The join is a boundary between two media along which the wave travels at different speeds, the wave being slower on the heavier rope. A wave meeting such a boundary cannot pass on undisturbed: the two ropes cannot keep their motions matched at the join unless some of the wave is sent back. Part of the energy is therefore carried forward into the heavy rope as a transmitted wave, and the rest is returned along the light rope as a reflected wave. Because it meets a heavier, slower medium, the reflected pulse is also inverted.

(b) Sound spreads out when it passes the edge of an obstacle or an opening, an effect called diffraction, and the spreading is large when the wavelength is comparable with, or larger than, the obstacle. The wavelength of ordinary speech is about a metre, comparable with the size of a wall or a doorway, so the sound diffracts strongly and bends round the corner to reach the listener. Light has a wavelength far smaller than any everyday obstacle, so it diffracts only negligibly and travels on in straight lines; hence the speaker cannot be seen.

(c) For simple harmonic motion the maximum speed is $v_{max} = \omega A$ and the maximum acceleration is $a_{max} = \omega^2 A$. Dividing one by the other eliminates the amplitude:

$$\omega = \frac{a_{max}}{v_{max}} = \frac{4.5m/s^2}{0.30m/s} = 15rad/s$$

The period follows from the angular frequency:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{15rad/s} = 0.42s$$

and the amplitude from $v_{max} = \omega A$:

$$A = \frac{v_{max}}{\omega} = \frac{0.30m/s}{15rad/s} = 0.020m$$

Question 24

(a) Take the displacement as $x = A\sin\omega t$. The velocity is the rate of change of displacement, $v = A\omega\cos\omega t$, and the acceleration is $a = -A\omega^2\sin\omega t$. The velocity is a quarter of a cycle, a phase angle of $\frac{\pi}{2}$, ahead of the displacement: it is greatest as the body crosses the centre, where the displacement is zero, and zero at the extremes, where the displacement is greatest. The acceleration is exactly half a cycle, a phase angle of π, out of step with the displacement: it is always directed opposite to the displacement, greatest at the extremes and zero at the centre.

(b) The ship sends a short pulse of sound vertically downwards. The pulse travels down to the sea-bed, is reflected there, and travels back up to a detector on the ship. The instrument measures the time between sending the pulse and receiving its echo. In that time the pulse has covered the depth twice, and since the speed of sound in seawater is known, the depth is half the product of that speed and the measured time. A longer delay therefore means deeper water.

(c) The speed of a transverse wave on a string is set by the tension and the linear mass density:

$$v = \sqrt{\frac{T}{\mu}} = \sqrt{\frac{50N}{2.0\times10^{-3}kg/m}} = 158m/s$$

The wavelength then follows from $v = f\lambda$:

$$\lambda = \frac{v}{f} = \frac{158m/s}{100Hz} = 1.58m$$

Question 25

(a) While the ambulance approaches, each successive compression of the sound wave is sent out from a point closer to the listener than the one before, so the compressions arrive more closely spaced: the effective wavelength is shortened and the frequency raised. While the ambulance recedes, each compression is sent out from a point further away, so they arrive more widely spaced: the wavelength is lengthened and the frequency lowered. At the instant of passing, the source changes from approaching to receding, so the pitch heard drops from the raised value to the lowered one.

(b) The stone gives the water a fixed amount of energy, and this energy is spread along the expanding circular ripple. As the ripple travels outwards its circumference grows, so the same energy must be shared along an ever-greater length of crest, and the energy carried by each portion of the ripple steadily decreases. Since the energy of a wave depends on the square of its amplitude, this falling energy shows itself as a falling amplitude, and the ripples grow weaker as they spread. Some energy is also lost in overcoming the viscosity of the water.

(c) For a mass on a spring $T = 2\pi\sqrt{m/k}$, which rearranges to give the force constant:

$$k = \frac{4\pi^2 m}{T^2} = \frac{4\pi^2(0.25\text{kg})}{(0.50\text{s})^2} = 39.5\text{N/m}$$

Since $T \propto \sqrt{m}$, doubling the period to 1.0s requires the mass to become four times as large, that is $4 \times 0.25\text{kg} = 1.0\text{kg}$.

Checking with the period formula:

$$m' = \frac{kT'^2}{4\pi^2} = \frac{(39.5\text{N/m})(1.0\text{s})^2}{4\pi^2} = 1.0\text{kg}$$

The further mass that must be added is therefore:

$$\Delta m = 1.0\text{kg} - 0.25\text{kg} = 0.75\text{kg}$$

Question 26

(a) Huygens' principle states that every point on a wavefront may be regarded as a source of secondary spherical wavelets, which spread out forwards with the speed of the wave, and that the new wavefront a short time later is the surface that just touches, the envelope of, all these secondary wavelets. To find how a wavefront advances, one draws from each point of the present wavefront a small wavelet, every wavelet of the same radius, equal to the wave speed multiplied by the time elapsed, and then draws the surface tangent to them all. That tangent surface is the wavefront in its new position.

(b) The metronome arm carrying its weight is an oscillating system that swings to and fro, and the rate of ticking is simply its natural frequency. That natural frequency depends on how the mass is distributed about the pivot. Sliding the weight outwards, away from the pivot, increases the rotational inertia of the arm, so each swing takes longer and the ticking slows. Sliding the weight inwards, towards the pivot, reduces the rotational inertia, so each swing is quicker and the ticking speeds up. The musician sets the tempo simply by choosing where the weight sits.

(c) The speed of sound in a gas is proportional to the square root of its absolute temperature, $v \propto \sqrt{T}$. Hence the speed at the higher temperature is:

$$v_2 = v_1\sqrt{\frac{T_2}{T_1}} = (340\text{m/s})\sqrt{\frac{308\text{K}}{288\text{K}}} = 352\text{m/s}$$

The wavelength of the 480Hz note follows from $\lambda = v/f$. At 15°C:

$$\lambda_1 = \frac{v_1}{f} = \frac{340\text{m/s}}{480\text{Hz}} = 0.708\text{m}$$

and at 35°C:

$$\lambda_2 = \frac{v_2}{f} = \frac{352\text{m/s}}{480\text{Hz}} = 0.733\text{m}$$

Question 27

(a) The pitch of a note, how high or low it sounds, is determined by the frequency of the sound wave: a higher frequency is heard as a higher pitch. The loudness of the note, how strong or faint it sounds, is determined by the intensity of the wave, and so by the amplitude of the vibration, since the intensity depends on the square of the amplitude: a larger amplitude is heard as a louder sound. The two properties are independent: two notes may share the same pitch but differ in loudness, or share the same loudness but differ in pitch, because each depends on a different feature of the wave.

(b) A person sitting on a swing is a simple pendulum, and the period of a simple pendulum, $T = 2\pi\sqrt{l/g}$, depends only on the length of the swing and on the acceleration of free fall, not on the mass of the person. A heavier person is pulled towards the lowest point by a larger gravitational restoring force, but that same person also has correspondingly greater inertia and so is harder to set moving; the two effects cancel exactly. The time for one complete to-and-fro swing is therefore the same for the heavy adult and the light child on a swing of the same length.

(c) The pulse travels to the far post and back, a total distance of $2 \times 100\text{m}$, at 40m/s:

$$t = \frac{2L}{v} = \frac{2(100\text{m})}{40\text{m/s}} = 5.0\text{s}$$

The speed of a wave on the cable is $v = \sqrt{T/\mu}$, so $v \propto \sqrt{T}$. Raising the tension to four times its value multiplies the speed by $\sqrt{4} = 2$, giving a new speed of 80m/s. The new return time is therefore:

$$t' = \frac{2L}{v'} = \frac{2(100\text{m})}{80\text{m/s}} = 2.5\text{s}$$

Question 28

(a) A wavefront is the surface reached by all the parts of the wave that are in step with one another. Close to a small source the wave has spread only a short distance, and the wavefront is a small circle, or sphere, centred on the source, with a pronounced curvature. As the wave travels outwards the radius of this circle grows; far from the source the radius is very large, and any small portion of so large a circle, which is the part an observer actually sees, is almost indistinguishable from a straight line. The wavefront therefore appears practically straight, a plane wavefront, at a great distance from the source.

(b) Ultrasound has a far higher frequency than audible sound, and therefore a much shorter wavelength. A wave can resolve, and so form a clear image of, detail only about as small as its own wavelength; the very short wavelength of ultrasound allows the fine structure of the baby to be made out, which long-wavelength audible sound could never show. Ultrasound is also strongly reflected at the boundaries between the different soft tissues, providing the echoes from which the image is built, and it can be sent in a narrow, well-directed beam. It is, in addition, safe for both mother and baby.

(c) The total energy of a simple harmonic oscillator is $E = \frac{1}{2}kA^2$, which gives the force constant:

$$k = \frac{2E}{A^2} = \frac{2(0.16\text{J})}{(0.080\text{m})^2} = 50\text{N/m}$$

At the centre of the motion all the energy is kinetic, $E = \frac{1}{2}mv_{max}^2$, which gives the maximum speed:

$$v_{max} = \sqrt{\frac{2E}{m}} = \sqrt{\frac{2(0.16\text{J})}{0.50\text{kg}}} = 0.80\text{m/s}$$

Question 29

(a) A driven system takes in energy from the driving force most effectively when it is pushed in step with its own natural motion. When the driving frequency matches the natural frequency, each push from the driving force acts in the same direction as the system's velocity, and so does positive work on it cycle after cycle, feeding energy in steadily. The amplitude grows until the rate at which energy is fed in is balanced by the rate at which it is lost to damping; when the damping is light this balance is reached only at a very large amplitude. At driving frequencies away from the natural frequency the force is repeatedly out of step, often opposing the motion, so far less energy is delivered and the amplitude stays small.

(b) Sound striking a surface is partly reflected and partly absorbed. Bare hard walls and floors reflect sound strongly and absorb very little, so a voice is reflected back and forth many times and dies away only slowly; the sound rings on and the room seems hollow. Soft furnishings, such as curtains, carpets and padded chairs, are good absorbers: their soft, porous surfaces convert the energy of the sound into heat at each encounter. Once the room is furnished the reflected sound is quickly absorbed, the ringing dies away rapidly, and the hollow quality disappears.

(c) In a stationary wave the distance from a node to the antinode next to it is one quarter of a wavelength, so:

$$\lambda = 4 \times 0.12\text{m} = 0.48\text{m}$$

The string vibrates 25 times each second, so $f = 25\text{Hz}$, and the speed of the waves follows from $v = f\lambda$:

$$v = f\lambda = (25\text{Hz})(0.48\text{m}) = 12\text{m/s}$$

Question 30

(a) The ear and brain cannot separate two sounds that arrive within about a tenth of a second of each other; such sounds merge into one. For a reflected sound to be heard as a distinct echo it must therefore arrive at least about a tenth of a second after the direct sound. The reflected sound has to travel from the listener to the surface and back again, so the surface must be far enough away for this extra journey to take more than about a tenth of a second; at the speed of sound in air this means a reflecting surface more than roughly 17m away. A nearer surface returns the sound too soon, and it is heard only as a prolonging of the original.

(b) The turned skipping rope is held at its two ends and carries a stationary wave. The held ends cannot move, so they must be nodes, points of zero displacement, while the middle is free to swing out fully and is an antinode. The simplest stationary wave that has a node at each end with one antinode between them is a single loop, half a wavelength long; this is the fundamental mode, the one set up at the steady, moderate rate at which the children naturally turn the rope. The rope therefore settles into one broad loop, its ends almost still and its middle sweeping out the widest path.

(c) For simple harmonic motion the speed at a displacement x is given by $v^2 = \omega^2(A^2 - x^2)$. Applying this at the two stated points:

$$(8.0\text{cm/s})^2 = \omega^2(A^2 - (6.0\text{cm})^2)$$
$$(6.0\text{cm/s})^2 = \omega^2(A^2 - (8.0\text{cm})^2)$$

Dividing the first equation by the second eliminates ω:

$$\frac{64}{36} = \frac{A^2 - 36}{A^2 - 64}$$

Cross-multiplying and solving for the amplitude:

$$64(A^2 - 64) = 36(A^2 - 36)$$
$$A = 10\text{cm}$$

Substituting back into the first equation gives the angular frequency:

$$\omega^2 = \frac{64}{A^2 - 36} = \frac{64}{64} = 1.0$$
$$\omega = 1.0\text{rad/s}$$

And hence the period:

$$T = \frac{2\pi}{\omega} = \frac{2\pi}{1.0\text{rad/s}} = 6.3\text{s}$$

Question 31

(a) In simple harmonic motion the speed is greatest at the centre and falls to zero at the extremes, since $v = \omega\sqrt{A^2 - x^2}$. Near the extremes the body is moving slowly, so it takes a relatively long time to cover any given small stretch of its path there; near the centre it is moving quickly and passes through in a short time. Over one complete cycle, therefore, more of the total time is spent in the slow-moving regions near the two extremes than in the fast-moving region near the centre.

(b) The breakwater has an end, and the open water beyond that end acts as a gap past which the sea waves travel. Waves passing the end of the wall are diffracted: they spread sideways into the region of calm water behind the wall. Because the wavelength of the sea waves is not small compared with the wall, this spreading is appreciable, so some of the wave energy curls round into the sheltered water and reaches the moored boats, setting them gently rocking. The waves there are weaker than in the open sea, since only the diffracted part of the wave arrives.

(c) The grating equation is $d\sin\theta = m\lambda$. For the first-order image, $m = 1$, the line spacing is:

$$d = \frac{m\lambda}{\sin\theta} = \frac{(1)(590 \times 10^{-9}\text{m})}{\sin 18.0^\circ} = 1.91 \times 10^{-6}\text{m}$$

This spacing is 1.91×10^{-3}mm, so the number of lines per millimetre of the grating is:

$$N = \frac{1\text{mm}}{1.91 \times 10^{-3}\text{mm}} = 524$$

Question 32

(a) At a dark fringe the light from the two slits arrives exactly out of step, the crest of one wave coinciding with the trough of the other. By the principle of superposition the two displacements cancel, so the resultant disturbance, and with it the intensity, is zero. The light energy is not destroyed but redistributed. Where the two waves arrive in step, at the bright fringes, they reinforce, and the intensity there is greater than the sum of the two separate intensities. Averaged over the whole pattern the total energy is exactly that supplied by the two slits; interference simply moves energy from the dark fringes to the bright ones.

(b) The clamped blade, with its free end, is an oscillating system; once twanged it vibrates freely at its own natural frequency, which is heard as the pitch of the buzz. That natural frequency depends on the stiffness of the vibrating part and on its mass. Sliding more of the blade out beyond the bench lengthens the vibrating portion: a longer projecting blade bends more easily, so it is less stiff, and it also carries more mass. Both changes lower the natural frequency, and so the buzz is heard to deepen.

(c) A string fixed at both ends and vibrating in its third harmonic carries three half-wavelengths along its length, so:

$$\lambda = \frac{2L}{3} = \frac{2(0.90\text{m})}{3} = 0.60\text{m}$$

The speed of the waves on the wire then follows from $v = f\lambda$:

$$v = f\lambda = (200\text{Hz})(0.60\text{m}) = 120\text{m/s}$$

Question 33

(a) A stationary wave is formed by two identical waves travelling through the medium in opposite directions. At a node the two waves always arrive exactly out of step: whatever displacement one wave would give to that point, the other gives an equal and opposite displacement at the same instant, and this holds at every instant. By the principle of superposition the two displacements cancel completely and permanently, so the medium at a node never moves. Each wave on its own would indeed move that point, but their contributions there are equal and opposite at all times, and their sum is always zero.

(b) Striking the fork harder gives its prongs more energy and so a larger amplitude of vibration, which is heard as a louder sound. But the frequency at which the prongs vibrate, and it is the frequency that fixes the musical note, is the natural frequency of the fork, set by the stiffness and the mass of the prongs. This natural frequency does not depend on the amplitude; the oscillation is isochronous. So however hard or gently the fork is struck, it vibrates at the same frequency and sounds the same note, differing only in loudness.

(c) The spacing of the bright fringes in a Young's pattern is $y = \frac{\lambda D}{d}$, which rearranges to give the wavelength:

$$\lambda = \frac{yd}{D} = \frac{(2.1\times10^{-3}\text{m})(0.40\times10^{-3}\text{m})}{1.5\text{m}}$$

which gives:

$$\lambda = 5.6\times10^{-7}\text{m}$$

Question 34

(a) A prism spreads light by refraction. Glass has a slightly greater refractive index for violet light than for red, so violet light is slowed and bent the more; violet is deviated most and red least. A grating spreads light by diffraction and interference: the direction of a bright order is fixed by $d\sin\theta = m\lambda$, so the angle of deviation increases with wavelength. Red light, having the longer wavelength, satisfies the grating equation at a larger angle than violet, so with a grating red is deviated most and violet least. The two instruments order the colours oppositely because they spread the light by entirely different mechanisms.

(b) Pushing and pulling the end of the coil squeezes the nearby turns together, forming a compression, and then draws them apart, forming a rarefaction. Each disturbed turn pushes and pulls on the next, so the compression is handed on from turn to turn and a longitudinal wave travels to the far end. Each individual turn of the coil only moves a short way back and forth about its own position; it is the pattern of compression, together with the energy carried with it, that travels the length of the coil.

(c) The force constant measures the force needed for each unit of extension. A given pull stretches the whole spring twice as far as it stretches either half alone, so each half is twice as stiff as the whole spring:

$$k' = 2k = 2(18\text{N/m}) = 36\text{N/m}$$

The period of the 0.40kg mass on one half-spring is then:

$$T = 2\pi\sqrt{\frac{m}{k'}} = 2\pi\sqrt{\frac{0.40\text{kg}}{36\text{N/m}}} = 0.66\text{s}$$

Question 35

(a) When a mass hangs on a vertical spring, gravity stretches the spring to a new equilibrium position at which the spring's pull exactly balances the weight. If the mass is then displaced from this position, the extra force that restores it is supplied entirely by the change in the spring's tension, just as it is for a horizontal spring displaced from its own equilibrium. The constant weight is already balanced and plays no further part in the motion. The restoring force per unit displacement is therefore the force constant of the spring in both cases, so the period $T = 2\pi\sqrt{m/k}$ is the same; gravity only shifts the equilibrium position, it does not change the period.

(b) The mist consists of a great many tiny water droplets. Light from the street lamp passing the droplets is diffracted, spreading slightly as it passes each small droplet, and the diffracted light from the many droplets interferes. The angle through which the light is bent depends on its wavelength, so each colour is reinforced at a slightly different angle away from the lamp. The eye therefore sees the lamp ringed by faint bands in which the colours are separated, a corona produced by diffraction.

(c) The wavelength follows from $v = f\lambda$:

$$\lambda = \frac{v}{f} = \frac{8.0\text{m/s}}{20\text{Hz}} = 0.40\text{m}$$

The number of complete waves contained in a 12m length is that length divided by one wavelength:

$$n = \frac{12\text{m}}{0.40\text{m}} = 30$$

Question 36

(a) As a wave passes, each particle of the medium is set oscillating, and the energy handed to a particle is the energy of its oscillation. The energy of an oscillating particle is proportional to the square of its amplitude: its greatest kinetic energy is $\frac{1}{2}mv_{max}^2$, and with $v_{max} = \omega A$ this is proportional to A^2. The wave passes this energy on from particle to particle, so the energy carried past any point each second, which is the power of the wave, is likewise proportional to the square of the amplitude. Doubling the amplitude therefore carries four times the energy, although the speed and frequency of the wave are unchanged.

(b) Light from each headlamp, on entering the eye, passes through the pupil, which is a small circular opening, and is diffracted there. Instead of forming a perfect point on the retina, each lamp forms a small spread-out patch of light. When the car is far away the two lamps are very close together in angle, so their two patches overlap and merge into what looks like a single light. As the car comes nearer the angle between the lamps grows; once it is large enough for the two patches to be seen apart, the lamps are made out as two separate lights.

(c) The total energy is $E = \frac{1}{2}kA^2$ and the potential energy at a displacement x is $\frac{1}{2}kx^2$, so $\frac{PE}{E} = \left(\frac{x}{A}\right)^2$. When the kinetic energy is three times the potential energy, the potential energy is one quarter of the total:

$$\left(\frac{x}{A}\right)^2 = \frac{1}{4}$$

$$x = \frac{A}{2} = \frac{0.10\text{m}}{2} = 0.050\text{m}$$

The kinetic energy is whatever part of the total is not potential, so the fraction that is kinetic is $1 - \left(\frac{x}{A}\right)^2$. At a displacement of 0.080m:

$$\frac{KE}{E} = 1 - \left(\frac{0.080\text{m}}{0.10\text{m}}\right)^2 = 0.36$$

Question 37

(a) The period of a mass on a spring is $T = 2\pi\sqrt{m/k}$. The period depends not on the mass itself but on its square root. Doubling the mass therefore multiplies the period not by 2 but by $\sqrt{2}$. The heavier mass does have more inertia and so is accelerated less by the same restoring force, which lengthens the period; but the period responds only to the square root of the mass, because it is $\sqrt{m}$, and not m, that appears in the formula. The period therefore rises by a factor of $\sqrt{2}$, about 1.4, and not by a factor of 2.

(b) A whip is not uniform: it tapers, growing thinner and lighter towards the tip. The speed of a transverse wave along it is $v = \sqrt{T/\mu}$, which is greater where the mass per unit length is smaller. As the wave runs towards the thin tip the mass per unit length steadily decreases, so the wave travels faster and faster. At the same time the energy carried by the wave is being delivered to less and less mass, so the amplitude of the motion grows, until the tip is flicked so violently that it produces the sharp crack.

(c) An ideal polaroid transmits one half of the intensity of unpolarised light, whatever its orientation, so after the first polaroid:

$$I_1 = \frac{I_0}{2} = \frac{16\text{W/m}^2}{2} = 8.0\text{W/m}^2$$

The light is now plane-polarised. By Malus's law the second polaroid, with its axis at 60° to the first, transmits:

$$I_2 = I_1(\cos 60°)^2 = (8.0\text{W/m}^2)(0.50)^2 = 2.0\text{W/m}^2$$

Question 38

(a) A bright fringe is formed wherever the path difference between the light from the two slits is a whole number of wavelengths: 0, λ, 2λ, 3λ and so on. In passing from one bright fringe to the next, the path difference increases by exactly one wavelength. For a screen far from the slits the path difference changes in direct proportion to the distance moved along the screen, so equal steps of one wavelength in path difference correspond to equal steps of distance on the screen. The bright fringes are therefore evenly spaced, the spacing being $\frac{\lambda D}{d}$.

(b) The child and the spring of the pogo stick together form a mass-on-a-spring system. When the spring is compressed below the equilibrium position it pushes back harder than the child's weight, and the restoring force is proportional to how far the spring is compressed. A restoring force proportional to the displacement is the condition for simple harmonic motion, whose period $T = 2\pi\sqrt{m/k}$ is fixed by the mass of the child and the stiffness of the spring, and does not depend on how hard the child bounces. Each bounce therefore takes the same time, giving the steady rhythm.

(c) A particle of the string performs simple harmonic motion, so its maximum speed is $v_{max} = \omega A$, with $\omega = 2\pi f$:

$$v_{max} = 2\pi f A = 2\pi(12\text{Hz})(0.015\text{m}) = 1.1\text{m/s}$$

The wave itself travels along the string at 9.0m/s. The maximum speed of a particle, 1.1m/s, is therefore much smaller than the speed of the wave; the two are quite distinct quantities.

Question 39

(a) When a train of waves crosses obliquely into shallower water it travels more slowly there. The wavefronts meet the boundary at an angle, so one end of a wavefront enters the shallow water and is slowed while the other end, still in deeper water, continues at the higher speed. Each wavefront therefore swings round as it crosses, and the whole train changes its direction of travel. The frequency is unchanged, so since $v = f\lambda$ the slower waves also have a shorter wavelength; the change of direction is the result of the two ends of each wavefront advancing unequally.

(b) When a wave tilts the boat to one side, the upward push of the water, the buoyancy force, and the weight of the boat no longer act along the same line, and together they produce a turning effect that acts to right the boat. This restoring turning effect grows with the angle of tilt, so once the wave has passed the boat rolls to and fro about its upright position in an oscillation that is close to simple harmonic. The rhythm of the roll is the boat's own natural period, fixed by its shape and by the way its mass is distributed, not by the wave that started the rolling.

(c) For Newton's rings the radius of the n-th dark ring is given by $r_n^2 = nR\lambda$, which rearranges to give the wavelength:

$$\lambda = \frac{r_n^2}{nR} = \frac{(2.4 \times 10^{-3}\text{m})^2}{(8)(1.2\text{m})}$$

which gives:

$$\lambda = 6.0 \times 10^{-7}\text{m}$$

Question 40

(a) Sunlight passing through the atmosphere is scattered by the air molecules, which are very much smaller than the wavelength of light. For scattering by such small particles the amount of scattering is far greater for short wavelengths than for long ones, being inversely proportional to the fourth power of the wavelength (Rayleigh scattering). Blue light, having a shorter wavelength than red, is therefore scattered much more strongly. Light reaching the observer's eye from all directions of the sky is this scattered light, in which the blue is dominant, so the sky appears blue.

(b) As waves approach a shelving beach the water grows shallower towards the shore, and waves travel more slowly in shallower water. When a wavefront comes in at an angle, the end of it nearer the shore is in shallower water and moves more slowly than the end still in deeper water. The slow end is held back while the fast end advances, so the wavefront steadily swings round until it lies almost parallel to the shoreline. By the time the waves reach the beach they are travelling nearly straight in, whatever angle they had far out; this turning is the refraction of the waves.

(c) In one complete oscillation the particle travels from one extreme to the other and back again, covering the amplitude four times, so the total path length is 4A:

$$A = \frac{24\text{cm}}{4} = 6.0\text{cm}$$

The angular frequency is $\omega = 2\pi/T$, and the maximum speed is $v_{max} = \omega A$:

$$\omega = \frac{2\pi}{T} = \frac{2\pi}{0.40\text{s}} = 15.7\text{rad/s}$$

$$v_{max} = \omega A = (15.7\text{rad/s})(0.060\text{m}) = 0.94\text{m/s}$$

and the maximum acceleration is $a_{max} = \omega^2 A$:

$$a_{max} = \omega^2 A = (15.7\text{rad/s})^2(0.060\text{m}) = 14.8\text{m/s}^2$$

PART THREE

Question 1

(a) The plates carry equal and opposite charges spread evenly over their facing surfaces, so each small patch of one plate has charge directly opposite it on the other; every such pair contributes the same field, and these add to give a field of equal strength and direction at all points between the plates (except near the edges), which is what is meant by a uniform field. An isolated point charge, by contrast, sends its field outward in all directions, and as the distance grows the same field lines spread over an ever larger spherical surface, so the field weakens as the inverse square of the distance.

(b) The diaphragm and the fixed plate together form a capacitor holding a fixed charge. A sound wave moves the diaphragm in and out, changing the plate separation and hence the capacitance. Since the charge Q is kept constant while C varies, the voltage $V = \frac{Q}{C}$ across the plates rises and falls in step with the sound, producing an electrical signal that copies the sound wave.

(c) (i) The capacitance is:

$$C = \frac{\varepsilon_0 A}{d} = \frac{8.85 \times 10^{-12}\text{F m}^{-1} \times 0.015\text{m}^2}{2.0 \times 10^{-3}\text{m}} = 66.4\text{pF}$$

(ii) the charge stored is:

$$Q = CV = 66.4 \times 10^{-12}\text{F} \times 200\text{V} = 1.33 \times 10^{-8}\text{C} = 13.3\text{nC}$$

(iii) the field between the plates is:

$$E = \frac{V}{d} = \frac{200\text{V}}{2.0 \times 10^{-3}\text{m}} = 1.0 \times 10^5\text{V m}^{-1}$$

(iv) and the energy stored is:

$$U = \frac{1}{2}CV^2 = \frac{1}{2} \times 66.4 \times 10^{-12}\text{F} \times (200\text{V})^2 = 1.33 \times 10^{-6}\text{J}$$

The capacitance is 66.4pF, the charge 13.3nC, the field 1.0×10^5V m^{-1}, and the energy stored 1.33μJ.

Question 2

(a) As charge is added to a capacitor the voltage across it rises in proportion, starting from zero when it is empty and reaching V only as the last of the charge arrives. The early charge is therefore moved across a small potential difference and the later charge across a larger one, so the average potential difference during charging is only $\frac{1}{2}$V. The work done, and hence the energy stored, is the charge times this average voltage, $\frac{1}{2}$QV, and not QV.

(b) The negative rod repels the free electrons in the metal cap and stem down into the leaves, so both leaves acquire a negative charge; being like charges they repel and the leaves diverge. When the rod is removed the electrons spread back evenly over the conductor, the leaves lose their excess charge and fall together again. No charge is transferred, because the rod never touches the electroscope.

(c) (i) The charge is conserved when the capacitors are joined, so the common voltage is the total charge divided by the total capacitance:

$$V = \frac{C_1 V_1}{C_1 + C_2} = \frac{4.0\,\mu\text{F} \times 300\text{V}}{4.0\,\mu\text{F} + 2.0\,\mu\text{F}} = 200\text{V}$$

(ii) the final charges are then:

$$Q_1 = C_1 V = 4.0\,\mu\text{F} \times 200\text{V} = 800\,\mu\text{C}, \qquad Q_2 = C_2 V = 2.0\,\mu\text{F} \times 200\text{V} = 400\,\mu\text{C}$$

(iii) and the energy lost is the difference between the initial and final stored energies:

$$\Delta U = 1/2\,C_1 V_1^2 - 1/2\,(C_1 + C_2)V^2 = 0.18\text{J} - 0.12\text{J} = 0.06\text{J}$$

The common voltage is 200V, the final charges are 800μC and 400μC, and 0.06J is lost (as heat and radiation in the connecting wires).

Question 3

(a) The potential at a point is the work done in bringing unit positive charge to it from a chosen reference point. For a point charge the field, and so the force on a test charge, dies away to nothing only at an infinite distance, so infinity is the natural place where a charge is free of the influence and is taken to have zero potential. The charge itself cannot be used as the zero, because there the field is infinite and the potential would be unbounded.

(b) When the switch opens, the current it was carrying cannot stop at once, and the voltage across the contacts tends to rise sharply and strike a spark. A capacitor connected across the contacts gives this charge somewhere to go: it absorbs the sudden surge, so the voltage across the gap rises only slowly and stays below the value needed to spark. The contacts are thereby protected from burning away.

(c) (i) The time constant is:

$$\tau = RC = 10 \times 10^3\Omega \times 50 \times 10^{-6}\text{F} = 0.50\text{s}$$

(ii) the initial discharge current is:

$$I_0 = \frac{V}{R} = \frac{100\text{V}}{10 \times 10^3\Omega} = 10\text{mA}$$

(iii) after one time constant the charge falls to a fraction e^{-1} of the initial charge $Q_0 = CV = 5.0\text{mC}$:

$$Q = Q_0 e^{-1} = 5.0\text{mC} \times 0.368 = 1.84\text{mC}$$

(iv) and all the initial stored energy ends up as heat in the resistor:

$$W = \frac{1}{2}CV^2 = \frac{1}{2} \times 50 \times 10^{-6}\text{F} \times (100\text{V})^2 = 0.25\text{J}$$

The time constant is 0.50s, the initial current 10mA, the charge after one time constant 1.84mC, and the heat dissipated 0.25J.

Question 4

(a) Capacitance is defined as the ratio of the charge stored to the voltage produced, $C = \frac{Q}{V}$. If more charge is placed on the plates the voltage rises in exact proportion, so the ratio is unchanged; the capacitance measures only how much charge the capacitor holds per volt, which is fixed by the area of the plates, their separation and the dielectric between them. It is therefore a property of the capacitor's construction, not of how much it happens to be charged.

(b) On a humid day a thin film of water, slightly conducting because of dissolved ions, settles on the sphere and on nearby insulators. This damp film provides a path along which the charge slowly leaks away to earth and to the air. On a dry day the film is absent, the surfaces stay good insulators, and the charge is held for much longer.

(c) (i) The capacitance of an isolated sphere is:

$$C = 4\pi\varepsilon_0 R = 4\pi \times 8.85 \times 10^{-12}\,\text{F m}^{-1} \times 0.15\text{m} = 16.7\text{pF}$$

(ii) the surface field reaches the breakdown value when $E = \frac{V}{R}$, so the greatest potential is:

$$V_{max} = ER = 3.0 \times 10^6\,\text{V m}^{-1} \times 0.15\text{m} = 4.5 \times 10^5\text{V}$$

(iii) the corresponding charge and (iv) energy are:

$$Q = CV_{max} = 16.7 \times 10^{-12}\text{F} \times 4.5 \times 10^5\text{V} = 7.5\,\mu\text{C}$$

$$U = \frac{1}{2}CV_{max}^2 = \frac{1}{2} \times 16.7 \times 10^{-12}\text{F} \times (4.5 \times 10^5\text{V})^2 = 1.7\text{J}$$

The sphere has a capacitance of 16.7pF, can be raised to 450kV, holding 7.5μC and storing 1.7J before the air breaks down.

Question 5

(a) Electric field is a vector while electric potential is a scalar. At the midpoint the two equal positive charges push outward in exactly opposite directions with equal strength, so their field vectors cancel and the resultant field is zero. Their potentials, however, are scalars that simply add: each charge raises the potential at the midpoint by the same positive amount, so the total is twice that of one charge and is certainly not zero.

(b) Glass is a dielectric whose molecules become polarised in the field of the plates, lining up so that their induced surface charges partly cancel that field. For a fixed charge this lowers the voltage across the plates, and since $C = \frac{Q}{V}$ a smaller voltage for the same charge means a larger capacitance. The glass therefore raises the capacitance above its air value by the factor of its dielectric constant.

(c) (i) The capacitance in air, and then with the mica, are:

$$C_0 = \frac{\varepsilon_0 A}{d} = \frac{8.85 \times 10^{-12}\,\text{F m}^{-1} \times 0.020\text{m}^2}{1.0 \times 10^{-3}\text{m}} = 177\text{pF}, \qquad C = \varepsilon_r C_0 = 5.0 \times 177\text{pF} = 885\text{pF}$$

(ii) the charge is fixed at $Q = C_0 V_0 = 17.7\text{nC}$ (the supply is disconnected), so the new voltage is:

$$V = \frac{Q}{C} = \frac{17.7 \times 10^{-9}\text{C}}{885 \times 10^{-12}\text{F}} = 20\text{V}$$

(iii) the field falls from $E_0 = \frac{V_0}{d} = 1.0 \times 10^5\,\text{V m}^{-1}$ to $E = \frac{V}{d} = 2.0 \times 10^4\,\text{V m}^{-1}$, and

(iv) the stored energy changes from:

$$U_0 = 1/2\,C_0 V_0^2 = 0.885\,\mu\text{J} \;\text{ to }\; U = 1/2\,CV^2 = 0.177\,\mu\text{J}$$

so the capacitance rises from 177pF to 885pF, the voltage falls to 20V, the field falls fivefold, and the stored energy drops by 0.71μJ (the mica being pulled in does the work).

Question 6

(a) The free electrons in a conductor move until no sideways push remains on them. If the field met the surface at a slant it would have a component along the surface, which would drive the surface electrons sideways and they would keep moving. Equilibrium is reached only when this tangential component has vanished, that is when the field meets the surface exactly at right angles.

(b) A defibrillator must pour a large amount of energy into the heart in a few milliseconds, far faster than the charging supply could deliver it. A capacitor is charged slowly and then discharged all at once, releasing its stored energy $\frac{1}{2}CV^2$ in a single brief, powerful pulse of controlled size. Connecting the patient straight to the supply could neither give so sharp a pulse nor limit the energy so safely.

(c) (i) For capacitors in series the reciprocals add:

$$\frac{1}{C} = \frac{1}{2.0\mu\text{F}} + \frac{1}{3.0\mu\text{F}} + \frac{1}{6.0\mu\text{F}} = 1.0\,\mu\text{F}^{-1} \;\Rightarrow\; C = 1.0\,\mu\text{F}$$

(ii) the same charge sits on each capacitor in series:

$$Q = CV = 1.0\,\mu\text{F} \times 90\text{V} = 90\,\mu\text{C}$$

(iii) and the voltage across each follows from $V = \frac{Q}{C}$:

$$V_2 = \frac{90\,\mu\text{C}}{2.0\,\mu\text{F}} = 45\text{V}, \quad V_3 = \frac{90\,\mu\text{C}}{3.0\,\mu\text{F}} = 30\text{V}, \quad V_6 = \frac{90\,\mu\text{C}}{6.0\,\mu\text{F}} = 15\text{V}$$

The combined capacitance is 1.0μF, each carries 90μC, and the voltages are 45V, 30V and 15V, which add to the 90V supply.

Question 7

(a) The force between two charges is governed by Coulomb's law, $F = \frac{Q_1 Q_2}{4\pi\varepsilon r^2}$, in which the medium enters through its permittivity ε. Oil has a higher permittivity than air, because its molecules polarise and partly screen the two charges from each other. This screening weakens the force, so the same charges the same distance apart push or pull less strongly in oil than in air, smaller by the factor of the oil's relative permittivity.

(b) An electrolytic capacitor uses an extremely thin layer of oxide, formed electrically on a metal foil, as its dielectric. Because the capacitance grows as the plate separation shrinks, this very small spacing yields an enormous capacitance from a modest area of foil, far more than a practical air gap could give. The thinness of the oxide layer is what packs so much capacitance into so small a case.

(c) (i) The initial force follows from Coulomb's law:

$$F_1 = \frac{kQ_1 Q_2}{r^2} = \frac{8.99 \times 10^9 \times 6.0 \times 10^{-9} \times 2.0 \times 10^{-9}}{(0.04\text{m})^2} = 6.7 \times 10^{-5}\text{N}$$

(ii) joining the two identical spheres lets the total charge share equally between them:

$$q = \frac{Q_1 + Q_2}{2} = \frac{6.0\text{nC} + 2.0\text{nC}}{2} = 4.0\text{nC}$$

(iii) so the new force is:

$$F_2 = \frac{kq^2}{r^2} = \frac{8.99 \times 10^9 \times (4.0 \times 10^{-9})^2}{(0.04\text{m})^2} = 9.0 \times 10^{-5}\text{N}$$

The force rises from 67μN to 90μN, because sharing makes the charges equal and a given total charge gives the greatest force when split equally.

Question 8

(a) The plates of a charged capacitor carry opposite charges and so attract each other, and to pull them apart this attraction must be worked against. With the capacitor disconnected the charge Q stays fixed, but a larger separation lowers the capacitance, and since the stored energy is $\frac{Q^2}{2C}$ a smaller C means more stored energy. The work done in separating the plates is exactly this increase in the electrostatic energy held in the field.

(b) The casing is a conductor, and in equilibrium any external field makes its free electrons rearrange until the field they set up cancels the external field within the metal and the space it encloses. The interior is therefore left with essentially no field, whatever fields act outside, so the sensitive circuit within is shielded.
(c) (i) The original capacitance and charge are:

$$C_0 = \frac{\varepsilon_0 A}{d} = \frac{8.85\times10^{-12}\times0.010}{1.0\times10^{-3}} = 88.5\text{pF}, \qquad Q = C_0V_0 = 4.43\text{nC}$$

(ii) tripling the gap thirds the capacitance:

$$C' = \frac{\varepsilon_0 A}{d'} = \frac{8.85\times10^{-12}\times0.010}{3.0\times10^{-3}} = 29.5\text{pF}$$

(iii) the charge is unchanged, so the voltage rises:

$$V' = \frac{Q}{C'} = \frac{4.43\times10^{-9}\text{C}}{29.5\times10^{-12}\text{F}} = 150\text{V}$$

(iv) and the work done equals the gain in stored energy:

$$W = 1/2\,C'V'^2 - 1/2\,C_0V_0^2 = 0.332\,\mu\text{J} - 0.111\,\mu\text{J} = 0.22\,\mu\text{J}$$

The capacitance falls from 88.5pF to 29.5pF, the voltage rises to 150V, and 0.22μJ of work is done in separating the plates.

Question 9

(a) By Gauss's law the field at a distance r from the centre depends only on the charge enclosed within a sphere of that radius and on the symmetry. For a uniformly charged sphere all the charge lies inside any external spherical surface, and the spherical symmetry makes the field the same in every direction, so it is identical to the field of a single point charge of the same total amount placed at the centre. Outside the sphere, then, its size makes no difference to the field.
(b) The battery can supply only a limited current, so it gives up its energy slowly; the capacitor, once charged, can release its stored energy in a tiny fraction of a second. Since power is energy divided by time, the same energy delivered in a far shorter time means a far greater power, so the brief discharge through the bulb is intensely bright. The battery alone, spread over a long time, could never reach that peak brightness.
(c) (i) The energy stored is:

$$U = \frac{1}{2}CV^2 = \frac{1}{2}\times2.0\times10^{-6}\text{F}\times(200\text{V})^2 = 0.040\text{J}$$

(ii) if this is delivered in 2.0ms the average power is:

$$P = \frac{U}{t} = \frac{0.040\text{J}}{2.0\times10^{-3}\text{s}} = 20\text{W}$$

(iii) which is four times the 5.0W the battery can supply, which is why the flash is so bright for its brief instant.

Question 10

(a) When first connected the uncharged capacitor offers no opposing voltage, so a large charging current flows and charge piles onto the plates. As the charge grows the capacitor's voltage rises and opposes the supply, so the current falls; once the capacitor's voltage equals the supply voltage there is no net push, the current falls to zero, and no steady current flows thereafter. A capacitor thus passes a brief charging current but blocks a steady direct current.
(b) The dome of a Van de Graaff generator has a very small capacitance, so even at several hundred thousand volts it holds only a tiny charge and stores very little energy; a touch discharges it in one harmless tingle. The mains sits at a much lower voltage but can drive a large current continuously for as long as contact lasts, delivering far more energy to the body. It is the sustained current and energy, not the high voltage alone, that make the mains lethal.
(c) (i) The time constant is:

$$\tau = RC = 1.5\times10^{6}\Omega\times2.0\times10^{-6}\text{F} = 3.0\text{s}$$

(ii) the maximum (final) charge is:

$$Q_{max} = CV = 2.0\times10^{-6}\text{F}\times9.0\text{V} = 18\,\mu\text{C}$$

(iii) after $t = 3.0\text{s} = \tau$ the charge has risen to:

$$Q = Q_{max}(1-e^{-t/\tau}) = 18\,\mu\text{C}\times(1-e^{-1}) = 11.4\,\mu\text{C}$$

(iv) and the charging current there is:

$$I = \frac{V}{R}e^{-t/\tau} = \frac{9.0\text{V}}{1.5\times10^{6}\Omega}\times e^{-1} = 2.2\,\mu\text{A}$$

The time constant is 3.0s, the maximum charge 18μC, the charge after 3.0s is 11.4μC, and the current there is 2.2μA.

Question 11

(a) In a uniform field the force on the charged particle is constant in both size and direction, so it gives the particle a constant acceleration at right angles to its original motion while leaving the motion along the original direction unchanged. This is exactly the situation of a projectile, which moves steadily forward while gravity accelerates it steadily downward. The combination of uniform velocity in one direction and uniform acceleration at right angles to it traces out a parabola in both cases.
(b) The dielectric of a capacitor can withstand only a certain field before it breaks down and starts to conduct. If the applied voltage is raised too high, the field across the dielectric exceeds its dielectric strength, a spark punches through it, and the capacitor is permanently ruined. The marked working voltage is set safely below this breakdown value.
(c) (i) The field between the plates is:

$$E = \frac{V}{d} = \frac{200\text{V}}{1.0\times10^{-2}\text{m}} = 2.0\times10^{4}\text{V m}^{-1}$$

(ii) the vertical acceleration is:

$$a = \frac{eE}{m} = \frac{1.6\times10^{-19}\text{C}\times2.0\times10^{4}\text{V m}^{-1}}{9.11\times10^{-31}\text{kg}} = 3.5\times10^{15}\text{m s}^{-2}$$

(iii) the time spent between the plates is $t = \frac{L}{v} = \frac{4.0\times10^{-2}\text{m}}{4.0\times10^{7}\text{m s}^{-1}} = 1.0\times10^{-9}\text{s}$, so the deflection is:

$$y = 1/2\,at^2 = 1/2\times3.5\times10^{15}\text{m s}^{-2}\times(1.0\times10^{-9}\text{s})^2 = 1.8\times10^{-3}\text{m}$$

(iv) and the vertical velocity gained is:

$$v_y = at = 3.5\times10^{15}\text{m s}^{-2}\times1.0\times10^{-9}\text{s} = 3.5\times10^{6}\text{m s}^{-1}$$

The field is 2.0×10^{4}V m^{-1}, the acceleration 3.5×10^{15}m s^{-2}, the deflection 1.8mm, and the vertical velocity 3.5×10^{6}m s^{-1}.

Question 12

(a) By Gauss's law the field around a long straight charged wire is found from a cylindrical surface whose area, $2\pi rL$, grows in proportion to the radius r. The same enclosed charge is therefore spread over a surface that increases as r, so the field falls as 1/r. For a point charge the surface is a sphere whose area grows as r^2, so its field falls off faster, as $1/r^2$.

(b) The inner and outer conductors of a coaxial cable are separated by an insulator, so they behave as the two plates of a capacitor and store charge whenever a voltage exists between them. This capacitance must be charged and discharged every time the signal changes, which draws current; at high frequencies the changes are so rapid that this current noticeably weakens and distorts the signal, so such cables are designed to keep their capacitance low.

(c) (i) The capacitance of the cable is:

$$C = \frac{2\pi\varepsilon_0 L}{\ln(b/a)} = \frac{2\pi \times 8.85 \times 10^{-12}\,\mathrm{F\,m^{-1}} \times 50\mathrm{m}}{\ln(4.0/1.0)} = 2.0\mathrm{nF}$$

(ii) and the charge stored at 200V is:

$$Q = CV = 2.0 \times 10^{-9}\mathrm{F} \times 200\mathrm{V} = 4.0 \times 10^{-7}\mathrm{C} = 0.40\,\mu\mathrm{C}$$

The cable has a capacitance of 2.0nF and stores 0.40μC at 200V.

Question 13

(a) Although the charge sits on the plates, the stored energy can be located in the space between them where the field exists, because the energy is exactly accounted for by that field. The stronger the field, the more energy it holds per unit volume, the amount being $1/2\,\varepsilon_0 E^2$. It is therefore natural and self-consistent to regard the electric field itself as the seat of the stored energy.

(b) The droplets all carry the same sign of charge, so they repel one another and spread out into a fine, even mist instead of clumping together. The metal object is earthed, so the charged droplets are attracted to it and follow the field lines, which curve round to its far side, carrying paint even onto the hidden surfaces. The result is an even coat with very little waste.

(c) (i) The capacitance is:

$$C = \frac{\varepsilon_0 A}{d} = \frac{8.85 \times 10^{-12} \times 0.025}{5.0 \times 10^{-3}} = 44.3\mathrm{pF}$$

(ii) the energy stored is:

$$U = 1/2\,CV^2 = 1/2 \times 44.3 \times 10^{-12}\mathrm{F} \times (1.0 \times 10^3\mathrm{V})^2 = 22.1\,\mu\mathrm{J}$$

(iii) the field is $E = \frac{V}{d} = 2.0 \times 10^5\mathrm{V\,m^{-1}}$, so the energy density is:

$$u = 1/2\,\varepsilon_0 E^2 = 1/2 \times 8.85 \times 10^{-12} \times (2.0 \times 10^5\mathrm{V\,m^{-1}})^2 = 0.177\mathrm{J\,m^{-3}}$$

and the volume between the plates is $Ad = 0.025\mathrm{m}^2 \times 5.0 \times 10^{-3}\mathrm{m} = 1.25 \times 10^{-4}\mathrm{m}^3$, so:

$$u \times (Ad) = 0.177\mathrm{J\,m^{-3}} \times 1.25 \times 10^{-4}\mathrm{m}^3 = 22.1\,\mu\mathrm{J}$$

which is equal to the stored energy, confirming that the energy resides in the field.

Question 14

(a) In a series arrangement the plates between neighbouring capacitors are joined by a conductor that began neutral, so charge pushed onto one plate draws an equal and opposite charge onto the next; every capacitor in the chain therefore carries the same charge. In a parallel arrangement all the upper plates are joined to one terminal and all the lower plates to the other, so each capacitor has the full supply voltage across it, and they share the same voltage while generally carrying different charges.

(b) A water molecule is polar, having its positive and negative ends slightly apart. The charged comb attracts the opposite end of each molecule a little more strongly than it repels the near end, so the stream as a whole feels a net pull toward the comb and bends visibly toward it, even though the water carries no net charge.

(c) (i) The parallel pair combine, and then add in series with the 3.0μF:

$$C_p = 2.0\,\mu\mathrm{F} + 4.0\,\mu\mathrm{F} = 6.0\,\mu\mathrm{F}, \qquad \frac{1}{C} = \frac{1}{3.0} + \frac{1}{6.0} = \frac{1}{2.0}\,\mu\mathrm{F}^{-1} \;\Rightarrow\; C = 2.0\,\mu\mathrm{F}$$

(ii) the charge drawn from the supply is:

$$Q = CV = 2.0\,\mu\mathrm{F} \times 60\mathrm{V} = 120\,\mu\mathrm{C}$$

(iii) this charge sits on the series 3.0μF and on the parallel block, across which the voltage is $V_p = \frac{120\,\mu\mathrm{C}}{6.0\,\mu\mathrm{F}} = 20\mathrm{V}$, so:

$$Q_2 = C_2 V_p = 2.0\,\mu\mathrm{F} \times 20\mathrm{V} = 40\,\mu\mathrm{C}, \qquad Q_4 = C_4 V_p = 4.0\,\mu\mathrm{F} \times 20\mathrm{V} = 80\,\mu\mathrm{C}$$

The equivalent capacitance is 2.0μF, the supply delivers 120μC, and the 2.0μF and 4.0μF capacitors carry 40μC and 80μC at 20V.

Question 15

(a) A high dielectric constant is wanted because it multiplies the capacitance, so the capacitor holds more charge and stores more energy at a given voltage. A high dielectric strength is wanted because it sets the greatest field, and hence the greatest voltage, the capacitor can take before the dielectric breaks down. A good dielectric needs both: a high constant so that it stores much, and a high strength so that it can be worked at a high voltage without failing.

(b) In flight an aircraft picks up static charge from friction with the air, which would build up to a high voltage and discharge noisily through the radio aerials, interfering with communications. At the sharp tips of the static wicks the field becomes very intense, so the charge leaks quietly away into the air at the points before it can accumulate, keeping the aircraft near the potential of its surroundings.

(c) The potential energy is the sum over the three pairs, $U = \frac{k}{r}(q_1q_2 + q_1q_3 + q_2q_3)$, with $r = 0.10\mathrm{m}$ and the products, in units of $10^{-12}\mathrm{C}^2$, equal to $(2)(2) + (2)(-3) + (2)(-3) = -8$:

$$U = \frac{8.99 \times 10^9\mathrm{N\,m^2C^{-2}} \times (-8.0 \times 10^{-12}\mathrm{C}^2)}{0.10\mathrm{m}} = -0.72\mathrm{J}$$

The total electric potential energy is −0.72J, the negative sign showing that the configuration is bound (net attractive).

Question 16

(a) The excess charges on a conductor are free to move and they all repel one another, so they spread as far apart as they can, which means onto the outer surface. Any charge remaining within the body or on an inner surface would feel a net outward push from the others and would move outward too, so in equilibrium none is left inside; the whole of the charge resides on the outer surface.

(b) A supercapacitor uses electrodes of enormous effective surface area together with an extremely thin charge-separation layer, and since the capacitance grows both with area and with thinness of the gap, this packs a very large capacitance into a small volume. It can therefore store a

large charge and release it rapidly, which suits it to delivering short bursts of power and to bridging brief interruptions where a battery would respond too slowly.
(c) (i) The time constant is:

$$\tau = RC = 2.0 \times 10^3\Omega \times 500 \times 10^{-6}F = 1.0s$$

(ii) the initial discharge current is:

$$I_0 = \frac{V}{R} = \frac{80V}{2.0 \times 10^3\Omega} = 40mA$$

(iii) all the stored energy ends up as heat in the resistor:

$$W = 1/2\,CV^2 = 1/2 \times 500 \times 10^{-6}F \times (80V)^2 = 1.6J$$

(iv) and the time for the charge to halve is:

$$t_{1/2} = 0.693\,\tau = 0.693 \times 1.0s = 0.69s$$

The time constant is 1.0s, the initial current 40mA, the energy dissipated 1.6J, and the half-time 0.69s.

Question 17

(a) The law says: The electrostatic force between two point charges falls off as the inverse square of their separation, so that doubling the distance reduces the force to a quarter. Strong evidence for this is that the field, and hence the force, is exactly zero everywhere inside a hollow charged conductor; such complete cancellation occurs only if the law of force is precisely inverse-square, so this null result is a sensitive test that confirms it.
(b) A key on a capacitive keyboard sits above a small capacitor; pressing the key brings one conducting layer closer to another, or brings a conducting finger near, which changes the plate separation and so the capacitance. The electronics continually sense each key's capacitance and register a keystroke the moment it changes, so no mechanical contact is needed.
(c) (i) The field between the plates is:

$$E = \frac{V}{d} = \frac{2.0 \times 10^3V}{5.0 \times 10^{-2}m} = 4.0 \times 10^4 V\,m^{-1}$$

(ii) the horizontal force on the sphere is:

$$F = qE = 5.0 \times 10^{-9}C \times 4.0 \times 10^4 V\,m^{-1} = 2.0 \times 10^{-4}N$$

(iii) and the thread settles at the angle for which this force balances against the weight:

$$\tan\theta = \frac{F}{mg} = \frac{2.0 \times 10^{-4}N}{2.0 \times 10^{-4}kg \times 9.8m\,s^{-2}} = 0.102 \Rightarrow \theta = 5.8°$$

The field is 4.0×10^4V m^{-1}, the force 2.0×10^{-4}N, and the thread hangs at 5.8° to the vertical.

Question 18

(a) Take a small Gaussian pillbox with one flat face just inside the conductor, where the field is zero, and the other just outside the surface. All the flux therefore passes out through the outer face alone, while the charge enclosed is the surface density σ times the area of that face. Gauss's law, flux equals enclosed charge divided by ε_0, then gives $E \times \text{area} = \frac{\sigma \times \text{area}}{\varepsilon_0}$, so the field just outside the surface is $E = \frac{\sigma}{\varepsilon_0}$.
(b) A high-voltage capacitor is immersed in oil because oil is a dielectric of much greater dielectric strength than air and can withstand the intense field between the plates without breaking down. The oil also seeps into any tiny air gaps, where breakdown would otherwise begin, and carries heat away. The capacitor can therefore be worked safely at a far higher voltage than it could in air.
(c) (i) The field between the plates is:

$$E = \frac{\sigma}{\varepsilon_0} = \frac{2.0 \times 10^{-6}C\,m^{-2}}{8.85 \times 10^{-12}F\,m^{-1}} = 2.26 \times 10^5 V\,m^{-1}$$

(ii) the potential difference is:

$$V = Ed = 2.26 \times 10^5 V\,m^{-1} \times 2.0 \times 10^{-3}m = 452V$$

(iii) the capacitance is:

$$C = \frac{\varepsilon_0 A}{d} = \frac{8.85 \times 10^{-12} \times 0.010}{2.0 \times 10^{-3}} = 44.3pF$$

(iv) and the charge on each plate is:

$$Q = \sigma A = 2.0 \times 10^{-6}C\,m^{-2} \times 0.010m^2 = 2.0 \times 10^{-8}C = 20nC$$

The field is 2.26×10^5V m^{-1}, the voltage 452V, the capacitance 44.3pF, and the charge 20nC.

Question 19

(a) The charge a sphere holds at a given potential is proportional to its capacitance, which is itself proportional to its radius. When the two spheres are joined by a wire they must come to the same potential, and since the larger sphere has the greater capacitance it takes the greater share of the charge at that common potential. Most of the charge therefore flows onto the large sphere.
(b) At a sharp point or small radius the surface charge crowds together and the field becomes very intense, easily reaching the value at which the surrounding air breaks down and the charge leaks away. By making high-voltage terminals large and well rounded, the surface charge is spread out and the field is kept low, so the air does not break down and the charge, and hence the high voltage, is held without leakage or corona.
(c) (i) The capacitance of the first sphere and its initial charge are:

$$C_1 = 4\pi\varepsilon_0 R_1 = 4\pi \times 8.85 \times 10^{-12} \times 0.05 = 5.56pF, \qquad Q = C_1V_1 = 0.50\,\mu C$$

(ii) on connection the charge is shared at a common potential set by the combined capacitance $4\pi\varepsilon_0(R_1 + R_2) = 16.7pF$:

$$V = \frac{Q}{4\pi\varepsilon_0(R_1 + R_2)} = \frac{0.50 \times 10^{-6}C}{16.7 \times 10^{-12}F} = 3.0 \times 10^4 V$$

(iii) and the final charges are in proportion to the radii:

$$Q_1 = C_1V = 0.167\,\mu C, \qquad Q_2 = 4\pi\varepsilon_0 R_2\,V = 0.334\,\mu C$$

The first sphere starts with 0.50μC; after connection the common potential is 3.0×10^4V and the charges are 0.167μC and 0.334μC, twice as much on the sphere of twice the radius.

Question 20

(a) The energy stored in a capacitor is $1/2\,CV^2$, which depends on the square of the voltage. Doubling the voltage therefore multiplies the stored energy by $2^2 = 4$, and not merely by two, because raising the voltage both doubles the charge and doubles the potential difference through which it is held. The stored energy thus increases fourfold.

(b) In dry weather, pulling synthetic clothing over the body rubs unlike materials together and transfers charge, leaving the garment and the skin oppositely charged. The opposite charges attract, so the cloth clings to the body, and where the charge jumps the small gaps it makes tiny sparks that are heard as crackles. In humid weather the charge leaks away through the moist air, so the clinging and crackling disappear.

(c) (i) The energy stored in the capacitor is:

$$U = 1/2\,CV^2 = 1/2 \times 10 \times 10^{-6}\text{F} \times (100\text{V})^2 = 0.050\text{J}$$

(ii) of which 60% is converted to gravitational potential energy:

$$E_{useful} = 0.60 \times 0.050\text{J} = 0.030\text{J}$$

(iii) so the mass is raised through:

$$h = \frac{E_{useful}}{mg} = \frac{0.030\text{J}}{0.050\text{kg} \times 9.8\text{m s}^{-2}} = 0.061\text{m} = 6.1\text{cm}$$

The capacitor stores 0.050J, delivers 0.030J usefully, and raises the mass through about 6.1cm.

Question 21

(a) The power dissipated in a resistor is the product of the voltage across it and the current through it, P = VI. Because the resistor obeys Ohm's law, V = IR, this can be rewritten in two further forms: putting V = IR gives P = I²R, while putting I = V/R gives P = V²/R. All three are equivalent; P = I²R is the most convenient when the current and resistance are known, and P = V²/R when the voltage and resistance are known.

(b) The heat produced in a resistor depends on the square of the current, and the rms value of an alternating current is defined as exactly the steady direct current that gives the same average heating. A 12V (rms) supply therefore drives the same average heating power through the element as a 12V steady supply, so the kettle boils in the same time on either. The element responds only to the mean-square current, so it cannot tell the two apart.

(c) (i) The operating resistance is:

$$R = \frac{V^2}{P} = \frac{(240\text{V})^2}{60\text{W}} = 960\Omega$$

(ii) the rms current and (iii) the peak current are:

$$I_{rms} = \frac{P}{V} = \frac{60\text{W}}{240\text{V}} = 0.25\text{A}, \qquad I_0 = I_{rms}\sqrt{2} = 0.25\text{A} \times \sqrt{2} = 0.35\text{A}$$

(iv) the peak instantaneous power is twice the average power:

$$P_{peak} = V_0 I_0 = 2P_{av} = 2 \times 60\text{W} = 120\text{W}$$

The resistance is 960Ω, the rms current 0.25A, the peak current 0.35A, and the peak power 120W; the power delivered fluctuates at twice the mains frequency, that is at 100Hz.

Question 22

(a) In an alternating current the driving voltage reverses many times a second, so the force on the free electrons reverses with it. The electrons are therefore pushed first one way and then the other, oscillating back and forth about fixed average positions instead of advancing steadily along the wire. Over a complete cycle their net displacement is zero, so no electron ever travels from the source to the appliance.

(b) The connecting wires are thick and of low-resistance copper, so although they carry the full current their resistance is small and they dissipate little heat, since the heating is I²R. The lamp filament is a thin wire of high resistance carrying the same current, so far more power is dissipated in it; this concentrated heating drives the filament to white heat while the leads stay cool.

(c) (i) The peak drift velocity is:

$$v_0 = \frac{I_0}{neA} = \frac{4.0\text{A}}{8.5 \times 10^{28}\text{m}^{-3} \times 1.6 \times 10^{-19}\text{C} \times 2.0 \times 10^{-6}\text{m}^2} = 1.5 \times 10^{-4}\text{m s}^{-1}$$

(ii) since the velocity varies as $v_0 \sin\omega t$, the displacement oscillates with amplitude:

$$x_0 = \frac{v_0}{\omega} = \frac{1.5 \times 10^{-4}\text{m s}^{-1}}{2\pi \times 50\text{Hz}} = 4.7 \times 10^{-7}\text{m}$$

The peak drift velocity is 1.5×10^{-4}m s^{-1}, and the electrons oscillate to and fro with an amplitude of only about 4.7×10^{-7}m.

Question 23

(a) To charge a battery, current must be driven backwards through it, into its positive terminal, against its own EMF. This is possible only if the charging source has the greater EMF, so that the net driving voltage, the difference of the two EMFs, is positive and can push the charging current through the combined internal resistances. If the source EMF were equal to or less than the battery's, no charging current would flow.

(b) The mains supplies alternating current at a high voltage, but a laptop needs a low, steady direct voltage. The transformer first steps the voltage down to the value required; the rectifier then turns the alternating current into a one-way (direct) current; and the smoothing capacitor finally levels out the pulses of the rectified output into a nearly steady direct voltage. All three stages are needed to convert the mains into usable d.c.

(c) (i) The charging current is set by the excess of the supply EMF over the battery EMF, divided by the total resistance:

$$I = \frac{V_{dc} - E}{R + r} = \frac{14\text{V} - 12\text{V}}{1.2\Omega + 0.30\Omega} = 1.33\text{A}$$

(ii) the terminal voltage of the battery while charging is:

$$V = E + Ir = 12\text{V} + 1.33\text{A} \times 0.30\Omega = 12.4\text{V}$$

(iii) and the rate at which energy is stored chemically is:

$$P_{chem} = EI = 12\text{V} \times 1.33\text{A} = 16.0\text{W}$$

The charging current is 1.33A, the battery terminal voltage 12.4V, and chemical energy is stored at 16.0W.

Question 24

(a) To a direct current a coil presents only the resistance of its wire, because once the current is steady the coil's magnetic field is unchanging and induces no back-EMF. To an alternating current, however, the ever-changing current sets up a continually changing flux that induces a back-EMF opposing the change; this appears as inductive reactance. The total opposition, the impedance $Z = \sqrt{R^2 + X_L^2}$, is therefore much greater than the resistance alone.

(b) An ammeter is connected in series, so the whole current passes through it; it must therefore have a very low resistance, so that inserting it scarcely alters the current it is meant to measure. A voltmeter is connected in parallel across a component; it must have a very high resistance, so that it draws almost no current and does not disturb the voltage it is meant to read.

(c) (i) The impedance is:

$$Z = \frac{V}{I} = \frac{60V}{1.0A} = 60\Omega$$

(ii) the inductive reactance and (iii) the inductance follow:

$$X_L = \sqrt{Z^2 - R^2} = \sqrt{(60\Omega)^2 - (30\Omega)^2} = 52.0\Omega, \qquad L = \frac{X_L}{2\pi f} = \frac{52.0\Omega}{2\pi \times 50\text{Hz}} = 0.165\text{H}$$

(iv) and the power dissipated (only in the resistance) is:

$$P = I^2R = (1.0A)^2 \times 30\Omega = 30W$$

The impedance is 60Ω, the reactance 52.0Ω, the inductance 0.165H, and the power dissipated 30W.

Question 25

(a) The useful (true) power of an a.c. load is VI cos φ, so the current it needs is I = P/(V cos φ). An inductive load has a power factor cos φ less than one, so for the same true power and voltage it must draw a larger current than a purely resistive load, whose power factor is one. The extra current carries the reactive power that surges back and forth without doing useful work.

(b) The heat generated in a cable is I²R, so a cable carrying a large current would overheat unless its resistance were kept low. Since resistance falls as the cross-sectional area rises (R = ρL/A), thick cables are used for large currents to keep the resistance, and so the heating and the voltage drop, small; a thin wire is adequate for a lamp, which draws only a small current.

(c) (i) The true power and (ii) the apparent power are:

$$P = VI\cos\phi = 240V \times 6.0A \times 0.75 = 1080W, \qquad S = VI = 240V \times 6.0A = 1440VA$$

(iii) a purely resistive load of the same true power would draw:

$$I_R = \frac{P}{V} = \frac{1080W}{240V} = 4.5A$$

(iv) so the extra power wasted in the cable because of the larger current is:

$$\Delta P = (I^2 - I_R^2)R = (6.0^2 - 4.5^2)A^2 \times 0.40\Omega = 6.3W$$

The true power is 1080W, the apparent power 1440VA, the resistive load would draw only 4.5A, and the low power factor wastes an extra 6.3W in the cable.

Question 26

(a) A capacitor passes alternating current because its plates are charged and discharged repeatedly as the voltage alternates, so a current flows in the wires at every instant; but it blocks direct current, because once charged to a steady voltage no further charge moves. An inductor passes direct current freely, since a steady current meets only the small resistance of its wire; but it opposes alternating current, because the changing current induces a back-EMF (inductive reactance) that grows with frequency.

(b) Tungsten is used because it has a very high melting point and can be raised to white heat without melting. The filament is tightly coiled so that a long length, and hence a high resistance, is fitted into a small space, concentrating the heat. It is sealed in an inert gas or a vacuum so that the hot metal does not burn away by oxidising in the air.

(c) (i) The capacitive reactance is:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50\text{Hz} \times 10 \times 10^{-6}\text{F}} = 318\Omega$$

(ii) the impedance and (iii) the rms current are:

$$Z = \sqrt{R^2 + X_C^2} = \sqrt{(2000\Omega)^2 + (318\Omega)^2} = 2.03 \times 10^3\Omega, \qquad I_{rms} = \frac{V}{Z} = \frac{12V}{2.03 \times 10^3\Omega} = 5.9\text{mA}$$

(iv) on direct current the capacitor blocks the steady current once charged, so:

$$I_{dc} = 0$$

The reactance is 318Ω, the impedance 2.03kΩ, the rms current 5.9mA, and the direct current is zero.

Question 27

(a) The rms value is used because it fixes the average heating effect, on which the safe carrying of current and the rating of equipment depend; the peak value, reached only for an instant, would overstate the steady effect by a factor √2. Marking the mains as 240V and a fuse as 13A in rms terms therefore states the values that govern the real power and heating, which is what matters for safety and design.

(b) The cost of running an appliance is its power multiplied by the time for which it runs, charged at a fixed rate per kilowatt-hour. A heater converts electrical energy into heat at a high rate, often two or three kilowatts, while a television uses only a small fraction of this. For the same running time the heater therefore uses far more energy and costs far more, even on the same mains.

(c) (i) The rms current and (ii) the peak current are:

$$I_{rms} = \frac{V}{R} = \frac{240V}{24\Omega} = 10A, \qquad I_0 = I_{rms}\sqrt{2} = 10A \times \sqrt{2} = 14.1A$$

(iii) the fuse responds to the rms (heating) current of 10A, which is below its 13A rating, so it does not blow, even though the current momentarily peaks at 14.1A. (iv) The power and the energy used are:

$$P = VI_{rms} = 240V \times 10A = 2400W = 2.4kW$$

$$\text{Energy} = Pt = 2.4\text{kW} \times 2.0\text{h} = 4.8\text{kWh}, \qquad \text{Cost} = 4.8\text{kWh} \times 290\ \text{TZS kWh}^{-1} = 1392\ \text{TZS}$$

The rms current is 10A and the peak 14.1A; the fuse holds, and the heater uses 4.8kWh, costing 1392 TZS.

Question 28

(a) When resistors are joined in series there is only one path, so the same current must pass through each in turn, while the supply voltage divides among them in proportion to their resistances. When they are joined in parallel each is connected directly across the same two supply terminals, so each has the full supply voltage across it, while the total current divides among the branches. Hence series elements share the current and parallel elements share the voltage.

(b) The iron core is magnetised first one way and then the other as the alternating current reverses, and a magnetised iron bar changes its length very slightly (magnetostriction). The core therefore expands and contracts twice in every cycle of the mains, setting the air around it vibrating; because two such movements occur per cycle, the hum is at twice the mains frequency, that is 100Hz on a 50Hz supply.

(c) (i) The parallel resistance is:

$$R_p = \frac{R_1R_2}{R_1 + R_2} = \frac{100\Omega \times 200\Omega}{100\Omega + 200\Omega} = 66.7\Omega$$

(ii) the inductive reactance and (iii) the impedance are:

$$X_L = 2\pi fL = 2\pi \times 50\text{Hz} \times 0.30\text{H} = 94.2\Omega, \qquad Z = \sqrt{R_p^2 + X_L^2} = \sqrt{(66.7\Omega)^2 + (94.2\Omega)^2} = 115\Omega$$

(iv) and the current drawn from the supply is:

$$I = \frac{V}{Z} = \frac{240\text{V}}{115\Omega} = 2.1\text{A}$$

The parallel resistance is 66.7Ω, the reactance 94.2Ω, the impedance 115Ω, and the supply current 2.1A.

Question 29

(a) A current-carrying wire contains equal amounts of positive charge (the fixed lattice ions) and negative charge (the free electrons), so the two cancel and the wire is electrically neutral. A current is merely the drift of the electrons through this stationary positive background; as one electron moves on, another enters to take its place, so the total charge in any region stays zero. The movement of charge does not mean any accumulation of charge.

(b) A filament lamp receives its power in pulses, the current falling to zero a hundred times a second, but the filament has thermal inertia and cannot cool appreciably in the hundredth of a second between pulses, so it glows at a nearly steady temperature. The eye, too, cannot follow such rapid changes (persistence of vision), so the lamp appears to shine steadily and no flicker is seen.

(c) (i) The rms current is:

$$I = \frac{E}{R + r} = \frac{250\text{V}}{49\Omega + 1.0\Omega} = 5.0\text{A}$$

(ii) the terminal voltage and (iii) the power delivered to the heater are:

$$V = IR = 5.0\text{A} \times 49\Omega = 245\text{V}, \qquad P_R = I^2R = (5.0\text{A})^2 \times 49\Omega = 1225\text{W}$$

(iv) and the power wasted inside the generator is:

$$P_r = I^2r = (5.0\text{A})^2 \times 1.0\Omega = 25\text{W}$$

The current is 5.0A, the terminal voltage 245V, the heater receives 1225W, and 25W is wasted inside the generator.

Question 30

(a) First, an alternating voltage can be stepped up and down efficiently by a transformer, so it can be sent at very high voltage (and therefore low current and small I^2R loss) and then reduced to a safe value for the consumer; a direct voltage cannot be changed so simply. Second, alternating current is generated directly and cheaply by rotating-coil generators. These advantages make a.c. far more convenient for large-scale generation and distribution.

(b) The kilowatt-hour is the unit in which electrical energy is sold; one kilowatt-hour is the energy used by an appliance of power one kilowatt running for one hour. It is preferred to the joule because the joule is inconveniently small for household amounts of energy, whereas one kilowatt-hour, equal to 3.6×10^6J, is a practical size that gives manageable numbers on a bill.

(c) (i) The secondary voltage and (ii) the secondary current are:

$$V_s = \frac{V_p}{20} = \frac{240\text{V}}{20} = 12\text{V}, \qquad I_s = \frac{V_s}{R} = \frac{12\text{V}}{0.60\Omega} = 20\text{A}$$

(iii) the power delivered to the coil and (iv) the primary current are:

$$P = V_sI_s = 12\text{V} \times 20\text{A} = 240\text{W}, \qquad I_p = \frac{P}{V_p} = \frac{240\text{W}}{240\text{V}} = 1.0\text{A}$$

(v) and the energy used in 5.0 hours is:

$$\text{Energy} = Pt = 0.240\text{kW} \times 5.0\text{h} = 1.2\text{kWh}$$

The secondary gives 12V at 20A, delivering 240W to the coil; the mains supplies 1.0A, and 1.2kWh is used in 5.0 hours.

Question 31

(a) In a single series loop there is only one path for charge, and since charge cannot pile up at any point, the same number of coulombs passes every cross-section each second. The current is therefore the same at all points, however the resistances differ; it is the voltage that divides among the components, not the current.

(b) A hot-wire ammeter responds to the heating produced by the current, which depends only on the mean square of the current and not on how rapidly it alternates. Provided the rms value is unchanged, the heating, and hence the reading, stays the same whatever the frequency, so the meter reads alike at 50Hz and at any other frequency.

(c) (i) The combined resistance is:

$$R_p = \frac{R_1R_2}{R_1 + R_2} = \frac{60\Omega \times 30\Omega}{60\Omega + 30\Omega} = 20\Omega$$

(ii) the rms current and (iii) the peak current are:

$$I_{rms} = \frac{V}{R_p} = \frac{240\text{V}}{20\Omega} = 12\text{A}, \qquad I_0 = I_{rms}\sqrt{2} = 17.0\text{A}$$

(iv) and the total power dissipated is:

$$P = \frac{V^2}{R_p} = \frac{(240\text{V})^2}{20\Omega} = 2880\text{W}$$

The combined resistance is 20Ω, the rms current 12A, the peak current 17.0A, and the power 2880W.

Question 32

(a) In a pure inductor or capacitor the current is a quarter-cycle out of step with the voltage, so during each cycle the component takes energy from the supply for one quarter and returns all of it during the next; the average power is therefore zero. A resistor has its current in step with its voltage, so it takes energy at every instant and turns it irreversibly into heat, giving a positive average power.

(b) A car battery is built with very low internal resistance, using large plates of big area set close together, so that even a modest terminal voltage can drive the hundreds of amperes the starter needs, since $I = E/(R + r)$ is large when r is tiny. A torch battery has a high internal resistance, so it can supply only a small current; demanding a large current would collapse its terminal voltage and exhaust it quickly.

(c) By itself the pure inductor draws:

$$I = \frac{V}{X_L} = \frac{240\text{V}}{60\Omega} = 4.0\text{A}, \qquad P = 0 \text{ (wattless)}$$

In series with the 80Ω resistor the impedance, current and power become:

$$Z = \sqrt{R^2 + X_L^2} = \sqrt{(80\Omega)^2 + (60\Omega)^2} = 100\Omega, \qquad I = \frac{240V}{100\Omega} = 2.4A$$

$$P = I^2R = (2.4A)^2 \times 80\Omega = 461W$$

Alone the inductor passes 4.0A but dissipates no power; with the resistor the current falls to 2.4A and 461W is dissipated, all of it in the resistor.

Question 33

(a) A potentiometer balances the cell's EMF against a known potential drop along a wire, the balance being found when the galvanometer shows no deflection. At balance no current is drawn from the cell, so there is no internal 'lost volts' (Ir), and the reading gives the true EMF. A voltmeter must draw some current to deflect, so it always reads the terminal voltage, which is a little less than the EMF.

(b) A choke limits the current by its inductive reactance, which opposes the alternating current without dissipating energy, since a pure inductor takes no average power. A resistor used instead would waste much energy as heat (I^2R). The choke therefore controls the current while keeping the running cost low, which is why fluorescent fittings use one.

(c) (i) The reactance of the choke is:

$$X_L = 2\pi fL = 2\pi \times 50Hz \times 0.15H = 47.1\Omega$$

(ii) the impedance and (iii) the current are:

$$Z = \sqrt{R^2 + X_L^2} = \sqrt{(40\Omega)^2 + (47.1\Omega)^2} = 61.8\Omega, \qquad I = \frac{240V}{61.8\Omega} = 3.88A$$

(iv) the power now dissipated in the heater is:

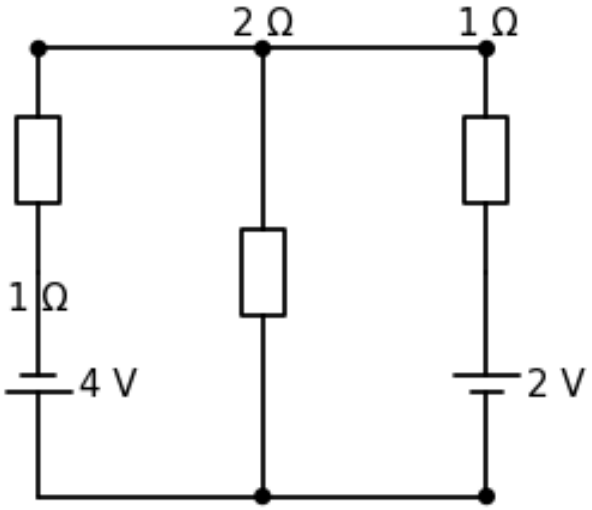
$$P = I^2R = (3.88A)^2 \times 40\Omega = 602W$$

Without the choke the current would be 240V/40Ω = 6.0A and the power 1440W, so the choke cuts the power from 1440W to about 602W while itself wasting almost none.

Question 34

(a) In a series a.c. circuit the instantaneous voltages are the actual voltages present at the same moment, and at every instant they must add up to the instantaneous supply voltage, since Kirchhoff's voltage law holds at each instant. Their rms or peak values, however, are reached at different instants because the components are out of phase, so these maxima never occur together and cannot simply be added; they must be combined as phasors instead.

(b) When two batteries of unequal EMF are joined in parallel, the one with the higher EMF tends to drive current round the loop, and if its EMF exceeds the other's sufficiently it forces current backwards through the weaker battery, charging it. Solving the network by Kirchhoff's laws then gives a negative value for the current in the weaker branch, showing that it flows opposite to the assumed direction.



(c) Let I_1 and I_2 be the currents from the 4.0V and 2.0V branches and I_3 the current in the central 2.0Ω resistor. Kirchhoff's current law at the top node gives $I_3 = I_1 + I_2$. Applying the voltage law to the left and right loops:

$$4.0V = (1.0\Omega)I_1 + (2.0\Omega)I_3, \qquad 2.0V = (1.0\Omega)I_2 + (2.0\Omega)I_3$$

Solving these three equations simultaneously:

$$I_1 = 1.6A, \qquad I_2 = -0.4A, \qquad I_3 = 1.2A$$

The 4.0V battery supplies 1.6A and the central resistor carries 1.2A; the negative sign for I_2 shows that 0.4A actually flows backwards through the 2.0V branch, which is being charged by the stronger battery.

Question 35

(a) On open circuit the secondary delivers no current, so the primary need draw only the small magnetising current that sets up the alternating flux in the core. This current is small because the primary winding has a large inductive reactance, which limits it; only enough flows to magnetise the core and supply the small iron losses. The transformer therefore takes very little power when nothing is connected to its secondary.

(b) The current to a distant load flows through the resistance of the long cable, which causes a voltage drop (IR) along it and wastes power as heat (I^2R). A thin cable has a large resistance, so the drop is big and the load receives a reduced voltage. Thick cables, of low resistance, are used so that the voltage drop and the power loss are kept small.

(c) (i) The current in the feeder is:

$$I = \frac{P}{V} = \frac{5000W}{250V} = 20A$$

(ii) the voltage drop and (iii) the voltage at the load are:

$$V_{drop} = IR = 20A \times 0.20\Omega = 4.0V, \qquad V_{load} = 250V - 4.0V = 246V$$

(iv) and the power wasted in the feeder is:

$$P_{loss} = I^2R = (20A)^2 \times 0.20\Omega = 80W$$

The feeder carries 20A, drops 4.0V, leaves 246V at the load, and wastes 80W.

Question 36

(a) A fuse protects a circuit by melting when the current becomes too large, and it melts because of the heating I^2R it suffers, which depends on the current it carries. Its safe rating is therefore set by that current, not by the voltage across it, which is almost zero while the fuse is intact. A fuse is accordingly marked with a current rating.

(b) The two carry the same current, but the voltage across the inductor leads the current by a quarter-cycle while the voltage across the capacitor lags it by a quarter-cycle. The two voltages are therefore half a cycle (180°) apart, exactly out of step: when one is at its positive peak the other is at its negative peak, so they oppose and tend to cancel.

(c) (i) The true power and (ii) the reactive power are:

$$P = VI\cos\phi = 240V \times 5.0A \times 0.80 = 960W, \qquad Q = VI\sin\phi = 240V \times 5.0A \times 0.60 = 720VAR$$

(iii) the impedance is $Z = \frac{V}{I} = \frac{240V}{5.0A} = 48\Omega$, so (iv) the resistance and reactance are:

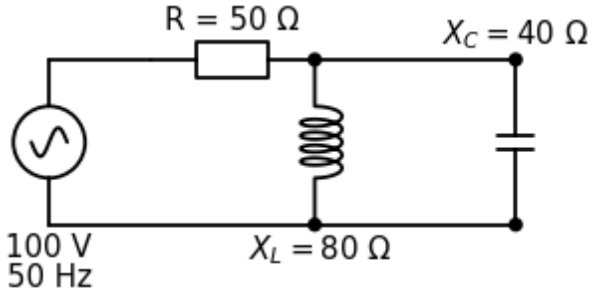
$$R = Z\cos\phi = 48\Omega \times 0.80 = 38.4\Omega, \qquad X = Z\sin\phi = 48\Omega \times 0.60 = 28.8\Omega$$

The true power is 960W, the reactive power 720VAR, the impedance 48Ω, the resistance 38.4Ω and the reactance 28.8Ω.

Question 37

(a) Impedance combines resistance with reactance, and the reactances depend on frequency: the inductive reactance $X_L = 2\pi fL$ rises with frequency while the capacitive reactance $X_C = \frac{1}{2\pi fC}$ falls. As the frequency changes, these reactances change, so the net reactance, and hence the impedance $Z = \sqrt{R^2 + (X_L - X_C)^2}$, also changes. The impedance of such a circuit is therefore a function of frequency.

(b) Switching on a high-power appliance draws a large current, which flows through the resistance of the house wiring and supply cable and causes a larger voltage drop (IR) there, so less voltage is left for the lamps and they dim. If the wiring is thin, and so of higher resistance, the effect is marked; thick, low-resistance wiring reduces it.



(c) (i) The inductor and capacitor in parallel present a net reactance:

$$X_{LC} = \frac{X_L X_C}{X_L - X_C} = \frac{80\Omega \times 40\Omega}{80\Omega - 40\Omega} = 80\Omega \text{ (capacitive)}$$

(ii) the impedance of the whole circuit and (iii) the supply current are:

$$Z = \sqrt{R^2 + X_{LC}^2} = \sqrt{(50\Omega)^2 + (80\Omega)^2} = 94.3\Omega, \qquad I = \frac{V}{Z} = \frac{100V}{94.3\Omega} = 1.06A$$

(iv) and the voltage across the parallel section is:

$$V_{LC} = IX_{LC} = 1.06A \times 80\Omega = 84.8V$$

The parallel section behaves as an 80Ω capacitive reactance; the supply delivers 1.06A, with 84.8V across the parallel combination.

Question 38

(a) Cells in series are joined so that the positive terminal of one connects to the negative of the next, so their EMFs all act in the same sense round the circuit and simply add. A chain of n equal cells therefore has a total EMF n times that of a single cell, able to drive current through a larger total resistance (their internal resistances add as well).

(b) The reactance of an inductor is $X_L = 2\pi fL$, proportional to the frequency. For direct current the frequency is zero, so the reactance is zero; once the current is steady the flux is unchanging and induces no back-EMF, and the inductor opposes the current only by the small resistance of its wire.

(c) (i) The secondary current is fixed by the rated apparent power:

$$I_s = \frac{S}{V_s} = \frac{2000VA}{24V} = 83.3A$$

(ii) the output power and (iii) the input power are:

$$P_{out} = S\cos\phi = 2000VA \times 0.90 = 1800W, \qquad P_{in} = P_{out} + 40W + 30W = 1870W$$

(iv) so the efficiency is:

$$\eta = \frac{P_{out}}{P_{in}} = \frac{1800W}{1870W} = 96.3\%$$

The secondary current is 83.3A, the output 1800W, the input 1870W, and the efficiency 96.3%.

Question 39

(a) An ideal transformer conserves power, so $V_pI_p = V_sI_s$. A step-down transformer has a lower secondary voltage, and since the power is unchanged the secondary current must be correspondingly larger. It therefore supplies a larger current than is drawn from the mains, at a lower voltage, which is why it is used where a large current at low voltage is needed.

(b) The energy a household uses is the power of its appliances multiplied by the time they run, and over a month this is an enormous number of joules. The kilowatt-hour, the energy used by a one-kilowatt appliance in one hour and equal to 3.6×10^6J, is a far more convenient size for such amounts, giving manageable figures, so the meter and the bill are reckoned in kilowatt-hours.

(c) (i) The inductive reactance is:

$$X_L = 2\pi fL = 2\pi \times 50Hz \times 0.30H = 94.2\Omega$$

(ii) the impedance and (iii) the current are:

$$Z = \sqrt{R^2 + X_L^2} = \sqrt{(80\Omega)^2 + (94.2\Omega)^2} = 123.6\Omega, \qquad I = \frac{240V}{123.6\Omega} = 1.94A$$

(iv) the voltages across the resistor and the coil are:

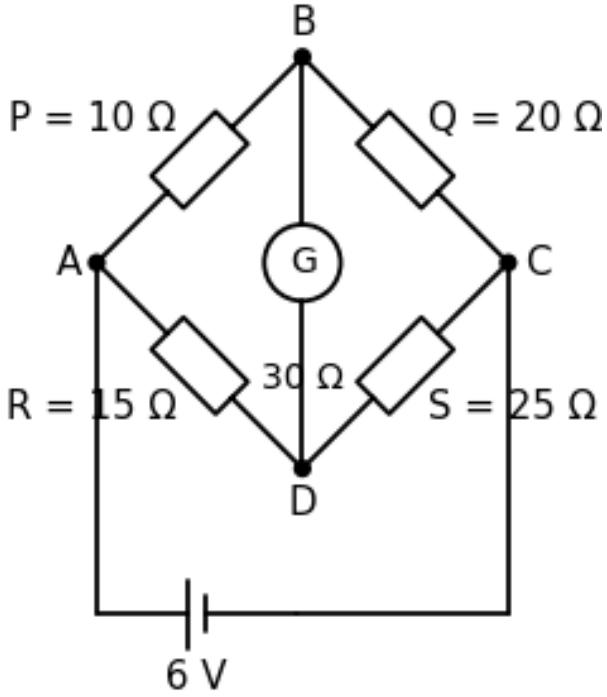
$$V_R = IR = 1.94A \times 80\Omega = 155.3V, \qquad V_L = IX_L = 1.94A \times 94.2\Omega = 183.0V$$

These add arithmetically to 338.3V, far more than the 240V supply; but because they are 90° out of phase they combine as a phasor sum, $\sqrt{V_R^2 + V_L^2} = \sqrt{(155.3V)^2 + (183.0V)^2} = 240V$, which equals the supply voltage.

Question 40

(a) A Wheatstone bridge is balanced when the ratio of the resistances in one pair of arms equals that in the other, $\frac{P}{Q} = \frac{R}{S}$, for then the two junctions are at the same potential and the galvanometer shows no deflection. At balance no current flows through the galvanometer, so the condition involves only the four arm resistances; the supply EMF affects the currents in the arms but not the ratio at which balance occurs, so the unknown resistance found is independent of the EMF.

(b) For a purely resistive load the current is in phase with the voltage, so the average power is $V_{rms}I_{rms}$, the same as a d.c. source of that voltage would deliver. For an inductive load the current lags the voltage, so only the part of it in phase with the voltage does work; the average power is $V_{rms}I_{rms}\cos\phi$, which is less, the reactive part of the current merely surging back and forth without delivering net power.



(c) The bridge is not balanced, since $\frac{P}{Q} = \frac{10}{20} = 0.5$ while $\frac{R}{S} = \frac{15}{25} = 0.6$, so a current flows through the galvanometer. With the galvanometer removed and A at 6.0V, C at 0, the potentials of B and D are:

$$V_B = E\frac{Q}{P+Q} = 6.0V \times \frac{20}{30} = 4.00V, \qquad V_D = E\frac{S}{R+S} = 6.0V \times \frac{25}{40} = 3.75V$$

so the open-circuit (Thévenin) voltage across BD is $V_{th} = V_B - V_D = 0.25V$.

The Thévenin resistance, with the battery short-circuited, is the two pairs of arms in parallel:

$$R_{th} = \frac{PQ}{P+Q} + \frac{RS}{R+S} = 6.67\Omega + 9.375\Omega = 16.0\Omega$$

and the galvanometer current is:

$$I_g = \frac{V_{th}}{R_{th} + R_g} = \frac{0.25V}{16.0\Omega + 30\Omega} = 5.4 \times 10^{-3}A = 5.4mA$$

The bridge is unbalanced and the galvanometer carries about 5.4mA.

Question 41

(a) The reactance of a capacitor is $X_C = \frac{1}{2\pi fC}$, which falls as the frequency rises. For a fixed voltage the current it passes is $\frac{V}{X_C}$, so a higher frequency, giving a smaller reactance, lets a larger current flow. In effect the plates charge and discharge more times each second, so more charge moves to and fro per second, which is a larger current.

(b) The figure 240V is the rms value; the voltage actually rises to a peak of $V_0 = V_{rms}\sqrt{2} \approx 339V$ at the crest of every cycle. The dielectric of the capacitor must withstand this peak without breaking down, so a working voltage safely above it, such as 400V, is chosen; a capacitor rated at only 240V would be over-stressed and could fail.

(c) (i) The reactance is:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50Hz \times 2.0 \times 10^{-6}F} = 1.59 \times 10^3\Omega$$

(ii) the rms current is:

$$I_{rms} = \frac{V}{X_C} = \frac{240V}{1.59 \times 10^3\Omega} = 0.151A$$

(iii) the maximum charge, at the peak voltage $V_0 = 240V \times \sqrt{2} = 339V$, is:

$$Q_{max} = CV_0 = 2.0 \times 10^{-6}F \times 339V = 6.79 \times 10^{-4}C = 679\,\mu C$$

(iv) and the maximum energy stored is:

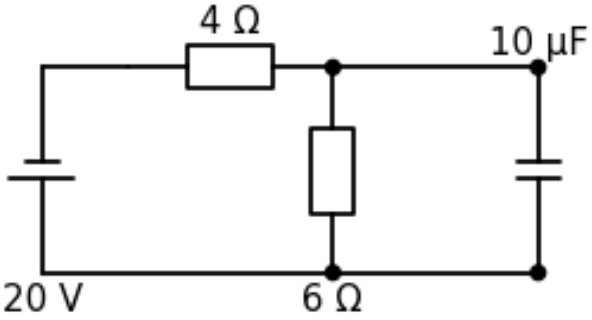
$$E_{max} = 1/2\,CV_0^2 = 1/2 \times 2.0 \times 10^{-6}F \times (339V)^2 = 0.115J$$

The reactance is 1.59kΩ, the rms current 0.151A, the maximum charge 679μC, and the maximum energy 0.115J.

Question 42

(a) Once a capacitor in a d.c. circuit is fully charged, no more charge moves onto its plates, so no current passes through its branch. Being connected in parallel with a component, it must share the same potential difference as that component; in the steady state it therefore simply settles at the voltage across whatever it is connected across.

(b) A large capacitor across the supply holds a reservoir of charge at the supply voltage. If the load suddenly draws more current the capacitor releases charge to help meet the surge, and if the demand falls it takes charge back up, so the terminal voltage is held steadier than the supply alone could manage.



(c) (i) In the steady state no current flows through the capacitor, so the current is that of the two resistors in series:

$$I = \frac{V}{R_1 + R_2} = \frac{20V}{4.0\Omega + 6.0\Omega} = 2.0A$$

(ii) the voltage across the 6.0Ω resistor is:

$$V_6 = IR_2 = 2.0A \times 6.0\Omega = 12V$$

(iii) the capacitor, in parallel with it, charges to this voltage, so:

$$Q = CV_6 = 10 \times 10^{-6}F \times 12V = 120\,\mu C$$

(iv) and the energy stored is:

$$E = 1/2\,CV_6^2 = 1/2 \times 10 \times 10^{-6}F \times (12V)^2 = 7.2 \times 10^{-4}J$$

The current is 2.0A, the voltage across the 6.0Ω resistor 12V, the charge on the capacitor 120μC, and the energy stored 0.72mJ.

Question 43

(a) Along a current-carrying wire the potential falls only gradually from one end to the other, so the field $E = \frac{V}{L}$ is small, and because the wire is uniform it is the same at every point. An isolated point charge, by contrast, concentrates its field close to itself, where it is very strong, and the field falls off as the inverse square of the distance; the two situations are quite different.

(b) The resistance of a wire is $R = \frac{\rho L}{A}$, so for the same length and material a thicker wire (larger A) has a smaller resistance. With the same voltage applied, the smaller resistance allows a larger current, since $I = \frac{V}{R}$, so the thick wire carries more current than the thin one.

(c) (i) The resistance is:

$$R = \frac{\rho L}{A} = \frac{1.7 \times 10^{-8}\Omega m \times 2.0m}{1.0 \times 10^{-6}m^2} = 0.034\Omega$$

(ii) the current and (iii) the field along the wire are:

$$I = \frac{V}{R} = \frac{1.5V}{0.034\Omega} = 44A, \qquad E = \frac{V}{L} = \frac{1.5V}{2.0m} = 0.75V\,m^{-1}$$

(iv) and the drift velocity of the electrons is:

$$v_d = \frac{I}{neA} = \frac{44A}{8.5 \times 10^{28}m^{-3} \times 1.6 \times 10^{-19}C \times 1.0 \times 10^{-6}m^2} = 3.2 \times 10^{-3}m\,s^{-1}$$

The resistance is 0.034Ω, the current 44A, the field 0.75V m^{-1}, and the drift velocity only 3.2×10^{-3}m s^{-1}.

Question 44

(a) An electric field exerts a force on a charge, and as the charge moves through a potential difference V the field does work qV on it. By the work-energy principle this work becomes kinetic energy, so $qV = 1/2\,mv^2$. The energy comes from the field, and ultimately from the source that maintains the potential difference, which gives up energy as the charge gains it.

(b) In a cathode-ray tube the electrons are accelerated through a potential difference and gain kinetic energy eV. A larger voltage does more work on each electron, so it reaches a higher speed, $v = \sqrt{\frac{2eV}{m}}$, before striking the screen; raising the accelerating voltage therefore produces a faster beam.

(c) (i) The kinetic energy gained is:

$$KE = eV = 1.6 \times 10^{-19}C \times 4.0 \times 10^{3}V = 6.4 \times 10^{-16}J = 4.0 \times 10^{3}eV$$

(ii) the final speed follows from $KE = 1/2\,mv^2$:

$$v = \sqrt{\frac{2\,KE}{m}} = \sqrt{\frac{2 \times 6.4 \times 10^{-16}J}{9.11 \times 10^{-31}kg}} = 3.7 \times 10^{7}m\,s^{-1}$$

(iii) the field in the gap and (iv) the force on the electron are:

$$E = \frac{V}{d} = \frac{4.0 \times 10^{3}V}{5.0 \times 10^{-2}m} = 8.0 \times 10^{4}V\,m^{-1}, \qquad F = eE = 1.6 \times 10^{-19}C \times 8.0 \times 10^{4}V\,m^{-1} = 1.3 \times 10^{-14}N$$

The electron gains 6.4×10^{-16}J (4.0keV), reaching 3.7×10^{7}m s^{-1}; the field is 8.0×10^{4}V m^{-1} and the force 1.3×10^{-14}N.

Question 45

(a) A full-wave rectifier turns both halves of each cycle the same way up, so its output is always positive; but it still rises and falls between zero and the peak a hundred times a second, so it is far from steady. To run electronic circuits, which need a nearly constant voltage, this pulsating output must be smoothed.

(b) A smoothing capacitor charges up to the peak and then feeds the load while the rectified voltage dips, so the output falls only a little before the next peak tops it up. A larger capacitor stores more charge for a given voltage ($Q = CV$), so it loses a smaller fraction of its charge, and hence of its voltage, during each dip, leaving a smaller ripple.

(c) (i) The mean output of a full-wave rectifier before smoothing is:

$$V_{dc} = \frac{2V_0}{\pi} = \frac{2 \times 12V}{\pi} = 7.64V$$

(ii) the charge on the reservoir capacitor at the peak and (iii) the energy stored are:

$$Q = CV_0 = 2200 \times 10^{-6}F \times 12V = 2.64 \times 10^{-2}C = 26.4mC, \qquad E = 1/2\,CV_0^2 = 1/2 \times 2200 \times 10^{-6}F \times (12V)^2 = 0.158J$$

(iv) and the mean load current before smoothing is:

$$I_{mean} = \frac{V_{dc}}{R} = \frac{7.64V}{60\Omega} = 0.127A$$

The mean output is 7.64V, the capacitor stores 26.4mC and 0.158J at the peak, and the mean load current is 127mA.

Question 46

(a) A steady current is a steady drift of the free electrons, and this drift is driven by the electric field set up inside the conductor by the potential difference across its ends. If the potential difference is removed, the field vanishes, the driving force disappears, and the electrons soon lose their net drift; a potential difference must therefore be maintained to keep a steady current flowing.

(b) Connecting wires are made of a good conductor, usually copper, and of generous thickness, so their resistance is extremely small compared with the components in the circuit. The voltage dropped across them is therefore negligible, and they are taken to have no resistance in order to simplify the analysis.

(c) (i) The current is:

$$I = \frac{E}{R + r} = \frac{12V}{5.5\Omega + 0.50\Omega} = 2.0A$$

(ii) the terminal voltage and (iii) the field along the 10cm resistor wire are:

$$V = IR = 2.0A \times 5.5\Omega = 11V, \qquad E = \frac{V}{L} = \frac{11V}{0.10m} = 110V\,m^{-1}$$

(iv) and the power dissipated in the resistor is:

$$P = I^2R = (2.0A)^2 \times 5.5\Omega = 22W$$

The current is 2.0A, the terminal voltage 11V, the field along the resistor 110V m^{-1}, and the power 22W.

Question 47

(a) A charged capacitor stores energy in the electric field between its plates, equal to $1/2\,CV^2$. When it is discharged this energy is released, and if it is released through a low resistance the charge flows out very quickly, giving a large current for a brief instant, enough to drive a device such as a small motor or a camera flash.

(b) The average current is the charge that flows divided by the time taken, and the average power is the energy delivered divided by the time. Releasing a given charge and a given energy in a shorter time therefore gives both a larger average current and a larger average power, which is why a rapid discharge is so vigorous.

(c) (i) The charge that flows and (ii) the energy released are:

$$Q = CV = 50 \times 10^{-6}F \times 100V = 5.0 \times 10^{-3}C = 5.0mC, \qquad E = 1/2\,CV^2 = 1/2 \times 50 \times 10^{-6}F \times (100V)^2 = 0.25J$$

(iii) the average current and (iv) the average power, over the 0.50s discharge, are:

$$I_{avg} = \frac{Q}{t} = \frac{5.0 \times 10^{-3}C}{0.50s} = 10mA, \qquad P_{avg} = \frac{E}{t} = \frac{0.25J}{0.50s} = 0.50W$$

The capacitor passes 5.0mC, releases 0.25J, and delivers an average of 10mA and 0.50W during the discharge.

Question 48

(a) An inductor and a capacitor in series resonate when their reactances are equal and opposite, $X_L = X_C$, for then they cancel and the circuit offers least opposition and responds most strongly. This happens at one particular natural frequency, $f_0 = \frac{1}{2\pi\sqrt{LC}}$, at which energy passes most freely back and forth between the two components.

(b) The capacitor stores energy in its electric field, $1/2\,CV^2$, and the inductor in its magnetic field, $1/2\,LI^2$. As the oscillation proceeds the current and charge are a quarter-cycle out of step, so when the charge is greatest (current zero) all the energy is in the capacitor, and a quarter-cycle later, when the current is greatest (charge zero), all of it is in the inductor.

(c) (i) The resonant frequency is:

$$f_0 = \frac{1}{2\pi\sqrt{LC}} = \frac{1}{2\pi\sqrt{0.40H \times 10 \times 10^{-6}F}} = 79.6Hz$$

(ii) at this frequency the two reactances are:

$$X_L = 2\pi f_0 L = 2\pi \times 79.6Hz \times 0.40H = 200\Omega, \qquad X_C = \frac{1}{2\pi f_0 C} = \frac{1}{2\pi \times 79.6Hz \times 10 \times 10^{-6}F} = 200\Omega$$

The resonant frequency is 79.6Hz, where the inductive and capacitive reactances are equal at 200Ω, confirming resonance.

Question 49

(a) When two capacitors are joined in series the same charge is pushed onto each, while the supply voltage divides between them; the combined capacitance is found from $\frac{1}{C} = \frac{1}{C_1} + \frac{1}{C_2}$. Because the reciprocals add, the combined capacitance is always smaller than either capacitor on its own, much as adding resistors in parallel gives a smaller resistance.

(b) While the capacitors are charging a current flows from the battery, but once they are fully charged no steady current flows. With no current there is no voltage dropped across the battery's internal resistance or any series resistor, so the whole of the battery's EMF appears across the capacitor combination.

(c) (i) The combined capacitance is:

$$C = \frac{C_1 C_2}{C_1 + C_2} = \frac{4.0\,\mu F \times 2.0\,\mu F}{4.0\,\mu F + 2.0\,\mu F} = 1.33\,\mu F$$

(ii) the same charge sits on each capacitor in series:

$$Q = CV = 1.33 \times 10^{-6} F \times 30V = 40\,\mu C$$

(iii) and the voltage across each follows from $V = \frac{Q}{C}$:

$$V_4 = \frac{40\,\mu C}{4.0\,\mu F} = 10V, \qquad V_2 = \frac{40\,\mu C}{2.0\,\mu F} = 20V$$

The combined capacitance is 1.33μF, each capacitor carries 40μC, and the voltages are 10V and 20V, which add to the 30V supply.

Question 50

(a) A capacitor blocks a steady direct current because, once it has charged to the supply voltage, no further charge can cross the gap between its plates and the current stops. It passes an alternating current because the continually reversing voltage keeps charging and discharging the plates, so charge flows back and forth in the connecting wires the whole time.

(b) On a.c. the capacitor opposes the current by its reactance $X_C = \frac{1}{2\pi fC}$, and a current flows; but the current leads the voltage by a quarter-cycle, so during each cycle the capacitor takes energy from the supply for one quarter and returns all of it during the next. The average power is therefore zero: the capacitor stores and returns energy but consumes none.

(c) (i) On the 100V battery the capacitor charges to:

$$Q = CV = 20 \times 10^{-6} F \times 100V = 2.0 \times 10^{-3} C = 2.0mC$$

after which the steady current and the power are both zero. (ii) On the 100V (rms), 50Hz supply the reactance, rms current and average power are:

$$X_C = \frac{1}{2\pi fC} = \frac{1}{2\pi \times 50Hz \times 20 \times 10^{-6} F} = 159\Omega, \qquad I_{rms} = \frac{V}{X_C} = \frac{100V}{159\Omega} = 0.628A$$

with an average power of zero, since the current and voltage are a quarter-cycle out of phase. On d.c. the capacitor holds 2.0mC and passes no current; on a.c. it draws 0.628A but consumes no power.

INDEX

A

Absolute zero, 76, 182, 187
Absorption spectrum, 778–780
Acceptor circuit. *See Resonance, series*
Adiabatic process, 243–252
  equation of, 245
  work done in, 248–250
  compared with isothermal, 251
Air wedge, 434–441
Alternating current (a.c.), 795–843
  defined, 796
  equation of, 796–800
  generation of, 796–800
  mean value of, 801–802
  measurement of, 833–836
  root-mean-square value of, 802–805
  why preferred for transmission, 826–829
Ammeter, 739–741
  conversion of galvanometer to, 739–741
  low resistance of, 739
Amplitude
  in simple harmonic motion, 285, 292–293
  of a wave, 367, 372
Ampère (unit), 702
Antinode, 385–388
Avogadro constant, 181, 185

B

Back-e.m.f., 798
Bandwidth (resonance), 823
Beats, 381–384
Bimetallic strip, 41, 781
Black-body radiation, 133–143
  emissivity, 134
  Stefan–Boltzmann law, 134
  Wien's displacement law, 135
Boltzmann constant, 185–187
Boyle's law, 180–181
Brewster's law (angle), 463–466
Brittleness, 18–20
Bulk modulus, 25–30

C

Capacitance, 641–645
  defined, 641
  of an isolated conductor, 684–689
  of a sphere, 684–686
  units of, 641
Capacitor, 641–700
  charging and discharging, 676–684
  combinations of (series and parallel), 651–657
  electrolytic, 643
  energy stored in, 657–665
  in a defibrillator, 682
  microphone, 847
  parallel-plate, 645–651
  smoothing (reservoir), 835
  supercapacitor, 642, 687
  working voltage of, 673
Carnot cycle and engine, 258–260
Cathode-ray tube, 619–620
Cell(s)
  e.m.f. and internal resistance, 725–732
  in series and in parallel, 733–735
  of unequal e.m.f. in parallel, 747–749
Celsius scale, 64, 78–80
Charge, electric, 553–558
  conservation of, 555
  coulomb (unit), 553
  quantisation of, 554
Charles's law, 180–181
Choke (inductor), 814, 833, 837
Coaxial cable, 687–688
Coefficient of linear expansion, 38–41
Coherence, 424–428
Compound (physical) pendulum, 323–329
Condenser. *See Capacitor*
Conduction, electrical
  in gases, 772–778
  in metals (drift of electrons), 702–708
Conduction, thermal, 109–127
Conductivity, electrical, 713–716
Continuous spectrum, 779
Convection, thermal, 127–133
Cooling, Newton's law of, 143–150
Coulomb (unit), 553
Coulomb's law, 558–565
  effect of medium (permittivity), 558, 664
  superposition principle, 558–564

Critical damping, 337–340

D

Damped oscillations, 336–341
Decibel scale, 497–503
Defibrillator, 682, 689
Degrees of freedom, 189–193
Dielectric, 664–675
  constant (relative permittivity), 664
  strength, 673
Diffraction, 458–467
Diffraction grating, 450–457, 779
  grating equation, 779
  in the spectrometer, 779
Diode, 833–836
Doppler effect, 513–531
  for light, and redshift, 524–529
  for sound, 513–524
  Hubble's law, 527–529
Drift velocity, 702–708
Ductility, 18–20

E

Eddy currents, 827, 837
  in a transformer core (loss), 827
  in induction heating, 837
Efficiency
  of a heat engine, 258–260
  of a transformer, 827–829
  at maximum power transfer, 766–769
Elastic limit, 16–17
Elastic potential energy, 30–34
Elasticity, 2–59
Electric current, 702–708
Electric field, 565–577
  from continuous distributions, 570–577
  from Gauss's law, 577–596
  motion of charge in uniform, 618–623
  of a point charge, 565–570
  relation to potential (E = −dV/dr), 610–618
Electric potential, 596–610
  equipotential surfaces, 610–618
  potential energy, 600–605
  zero at infinity, 596
Electrolysis, Faraday's laws of, 782
Electromagnetic spectrum, 510–513
Electromagnetic waves, 503–510
Electromotive force (e.m.f.), 725–732
Electron-volt, 600
Electroscope, gold-leaf, 555
Electrostatics, 552–639
Energy
  in simple harmonic motion, 298–302
  internal, 189–193, 213–217
  stored in a capacitor, 657–665
  density of electric field, 661
Entropy, 257–260
Equipotential surfaces, 610–618
Expansion, thermal, 38–44
  linear, area and volume, 38–41
  applications, 44

F

Farad (unit), 641
Faraday's law of induction, 796, 826
First law of thermodynamics, 213–218
Floating body, oscillation of, 316–319
Fluorescent lamp, 781–782, 814, 836
Forced oscillations, 341–345
Form factor, 803, 835
Fraunhofer lines, 450
Free oscillations, 341
Fundamental frequency, 393–397
Fuse, 768, 781

G

Galvanometer, 739–741
  conversion to ammeter and voltmeter, 739–741
  figure of merit, 786
Gas discharge, 772–778
  breakdown potential, 773
  initiated and self-sustaining, 776–777
Gauss's law, 577–596
  four canonical applications, 577–595
Generator, a.c., 796–800

H

Harmonics
  in pipes (air columns), 490–497
  on strings, 393–397
Heat
  capacity. *See Specific heat capacity*
  engine, 258–260
  latent, 88–94
  transfer, 107–168
Henry (unit), 809
Hooke's law, 7–10, 302

Hubble's law, 527–529
Huygens' principle, 418–424
Hysteresis
  elastic, 21–22
  magnetic (transformer loss), 827

I

Ideal gas, 180–185
  assumptions of kinetic theory, 170–173
  equation, 180–185
  real-gas limits, 193–194
Impedance, 813–818
  in a parallel circuit, 818–822
Inductance, 809
Induction cooker, 837
Inductor, in a.c. circuit, 808–813
Intensity
  of sound (decibel scale), 497–503
  of a wave, 372–376
Interference, 424–458
  conditions for, 424–428
  in thin films and air wedges, 434–441
  Newton's rings, 441–458
  versus diffraction, 458–467
  Young's double slit, 428–434
Internal energy, 189–193, 213–217
Internal resistance, 725–732
Inverse-square law, 558, 565
Isobaric process, 234–239
Isochoric process, 230–234
Isothermal process, 239–243

J

Joule heating, 764–772
Joule (unit), 209
Joule's law of heating, 765

K

Kelvin (thermodynamic) scale, 76–81, 182, 187
Kinetic theory of gases, 169–207
Kirchhoff's laws, 745–754
  current law, 745
  voltage law, 746
  node-voltage method, 752–753

L

Laplace correction (speed of sound), 486–488
Latent heat, 88–94
  of fusion and vapourisation, 88–92
LC oscillations, 822–826
Lenz's law, 798
Light, speed of, 503
Lightning conductor, 623–624
Line spectrum, 779
Longitudinal waves, 361–365, 480
Loudness, 497–503
Loudspeaker, 349, 837

M

Malus's law, 462
Mass–spring system, 302–309
Maximum power transfer theorem, 766–769
Mean value of alternating current, 801–802
Metre bridge, 754–759
Millikan's oil-drop experiment, 621–622
Modulus
  bulk, 25–30
  shear, 25–28
  Young's, 10–16
Molar heat capacities of a gas, 223–230
Moment of inertia, 323–329

N

Newton's law of cooling, 143–150
Newton's rings, 441–458
Nichrome, 767
Node, 385–388

O

Ohm (unit), 709
Ohm's law, 708–713
  microscopic derivation, 708–712
  ohmic and non-ohmic conductors, 784
Optical spectra of gases, 778–781
Oscillations. *See Simple harmonic motion*

P

Parallel-plate capacitor, 645–651
Pendulum
  compound (physical), 323–329
  simple, 309–316
Permittivity, 558, 664
Phase and phase difference, 372–376
Phasor, 805–808
Pipes, resonance in (open and closed), 490–497
Plastic deformation, 17–20
Polarisation, 460–467
  by reflection (Brewster's angle), 463–466
  Malus's law, 462
Potential, electric. *See Electric potential*

Potential energy
elastic, 30–34
electric, 600–605
intermolecular curve, 3–7
Potentiometer, 759–764
comparison of e.m.f.s, 760
measurement of internal resistance, 761
Power
electrical, 764–766
factor, 831–833
factor correction, 833
in a.c. circuits (true, reactive, apparent), 829–833
maximum transfer of, 766–769
Pressure of an ideal gas, 173–180
Progressive (travelling) wave, 366–372
PV diagram, 252–257

Q

Q-factor (quality factor), 823
Quantisation of charge, 554

R

Radiation, thermal, 133–143
Reactance
capacitive, 810–812
inductive, 809–810
Real gases, 193–194
Rectification, 833–836
half-wave and full-wave, 834
smoothing capacitor, 835
Redshift, 524–529
Rejector (tank) circuit. *See Resonance, parallel*
Resistance, 709
internal, 725–732
temperature coefficient of, 716–722
Resistivity, 713–716
Resistors, in series and parallel, 732–735
Resonance
electrical, series and parallel, 822–826
in pipes, 490–497
in simple harmonic motion, 341–345
Rolling oscillation on a curved surface, 333–336
Root-mean-square (r.m.s.) speed of gas molecules, 176–177
Root-mean-square (r.m.s.) value of a.c., 802–805

S

Second law of thermodynamics, 257–261
Shear modulus, 25–28
Shunt, 739–741
Simple harmonic motion (S.H.M.), 284–359
defining condition, 287–292
energy in, 298–302
equations of, 292–298
vocabulary of oscillation, 285–287
Simple pendulum, 309–316
Snell's law, 420
Sonometer, 393–400
Sound, 479–533
intensity and the decibel, 497–503
nature of, 480–486
speed of, 486–490
Specific heat capacity, 83–88
Spectra, optical, 778–781
emission and absorption, 779
line, band and continuous, 779
Spectrometer, 779
Standing (stationary) waves, 384–393
equation of, 388–393
Stefan–Boltzmann law, 134
Strain, 7–10
Stress, 7–10
Stress–strain curve, 16–25
Supercapacitor, 642, 687
Superposition
of charges and fields, 558–564, 570
of waves, 376–384

T

Tank circuit. *See Resonance, parallel*
Temperature, 60–106
and molecular kinetic energy, 185–189
scales, 63–67, 76–81
Temperature coefficient of resistance, 716–722
Terminal voltage, 725–732
Thermal conductivity, 109–111, 150–154
Thermal expansion, 38–44
Thermal radiation, 133–143
Thermocouple, 71
Thermodynamics, 208–273
first law, 213–218
second law, 257–261
Thermometers, types of, 67–76
Thin films, 434–441
Time constant (RC), 678–684
Torsional oscillation, 329–333
Transformer, 826–829

efficiency of, 827–829
losses in (copper, hysteresis, eddy), 827
step-up and step-down, 827
turns ratio, 826
Transverse waves, 361–365
Triple point, 76–80

U

Uniform electric field, motion in, 618–623
U-tube, oscillation of liquid in, 316

V

Van de Graaff generator, 576, 623
Velocity
drift, 702–708
of mechanical waves, 400–403
of vibration of a string, 345–348
Vibrations of strings, 393–400
Voltage magnification (resonance), 823
Voltmeter, 739–741

W

Wave motion, 360–416
intensity, 372–376
longitudinal and transverse, 361–365
progressive wave equation, 366–372
speed in different media, 400–403
stationary waves, 384–393
superposition, 376–384
Wavefront, 418–424
Wheatstone bridge, 754–759
balanced, 735–736
unbalanced, 752–753
Wien's displacement law, 135
Work done by a gas, 218–223

X

X-rays, 512, 837

Y

Yield point, 17
Young's double-slit experiment, 428–434
Young's modulus, 10–16
experimental determination, 34–38

Z

Zeroth law of thermodynamics, 60–63