π	pi	3	
G	Newton's constant	$7 \cdot 10^{-11}$	$\mathrm{kg^{-1}\,m^3\,s^{-1}}$
c	speed of light	$3 \cdot 10^{8}$	$\mathrm{m\,s^{-1}}$
k_B	Boltzmann's constant	10^{-4}	$\mathrm{eV\,K^{-1}}$
e	electron charge	$1.6 \cdot 10^{-19}$	C
σ	Stefan–Boltzmann constant	$6 \cdot 10^{-8}$	$\mathrm{W\,m^{-2}\,K^{-4}}$
m_sun	Solar mass	$2 \cdot 10^{30}$	kg
R_earth	Earth radius	$6 \cdot 10^{6}$	m
$\theta_\mathrm{moon/sun}$	angular diameter	10^{-2}	
ρ_air	air density	1	$\mathrm{kg\,m^{-3}}$
ρ_rock	rock density	5	$\mathrm{g\,cm^{-3}}$
$\hbar c$		200	eV nm
$L_\mathrm{vap}^\mathrm{water}$	heat of vaporization	2	$\mathrm{MJ\,kg^{-1}}$
γ_water	surface tension of water	10^{-1}	$\mathrm{N\,m^{-1}}$
a_0	Bohr radius	0.5	Å
a	typical interatomic spacing	3	Å
N_A	Avogadro's number	$6 \cdot 10^{23}$	
$\mathcal{E}_\mathrm{fat}$	combustion energy density	9	$\mathrm{kcal\,g^{-1}}$
E_bond	typical bond energy	4	eV
$\dfrac{e^2/4\pi\epsilon_0}{\hbar c}$	fine-structure constant α	10^{-2}	
p_0	air pressure	10^{5}	Pa
ν_air	kinematic viscosity of air	$1.5 \cdot 10^{-5}$	$\mathrm{m^2\,s^{-1}}$
ν_water	kinematic viscosity of water	10^{-6}	$\mathrm{m^2\,s^{-1}}$
day		10^{5}	s
year		$\pi \cdot 10^{7}$	s
F	solar constant	1.3	$\mathrm{kW\,m^{-2}}$
AU	distance to sun	$1.5 \cdot 10^{11}$	m
P_basal	human basal metabolic rate	100	W
K_air	thermal conductivity of air	$2 \cdot 10^{-2}$	$\mathrm{W\,m^{-1}\,K^{-1}}$
K	… of non-metallic solids/liquids	1	$\mathrm{W\,m^{-1}\,K^{-1}}$
K_metal	… of metals	10^{2}	$\mathrm{W\,m^{-1}\,K^{-1}}$
$c_\mathrm{p}^\mathrm{air}$	specific heat of air	1	$\mathrm{J\,g^{-1}\,K^{-1}}$
c_p	… of solids/liquids	25	$\mathrm{J\,mole^{-1}\,K^{-1}}$

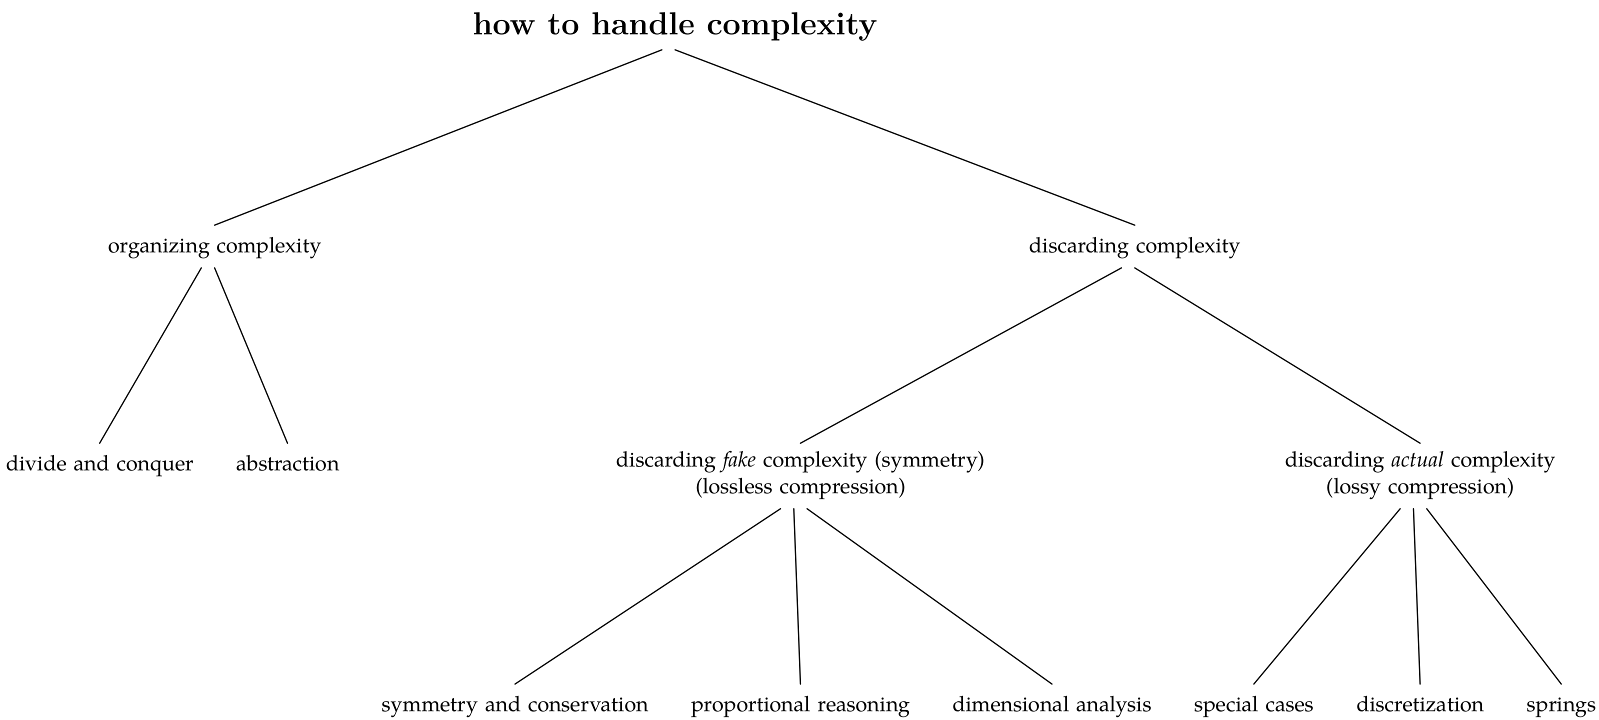
how to handle complexity
organizing complexity
discarding complexity
divide and conquer
abstraction
discarding fake complexity (symmetry)
(lossless compression)
discarding actual complexity
(lossy compression)
symmetry and conservation
proportional reasoning
dimensional analysis
special cases
discretization
springs

Contents

Preface

An approximate analysis is often more useful than an exact solution!

This counterintuitive thesis, the reason for this book, suggests two questions.

One question is: If science and engineering are about accuracy, how can approximate models be useful? They are useful because our minds are a small part of the world itself. When we represent a piece of the world in our minds, we discard many aspects – we make a model – in order that the model fit in our limited minds. An approximate model is all that we can understand. Making useful models means discarding less important information so that our minds may grasp the important features that remain.

This perhaps disappointing conclusion leads to a second question: Since every model is approximate, how do we choose useful approximations? The American psychologist William James said [10, p. 390]: 'The art of being wise is the art of knowing what to overlook.' This book therefore develops intelligence amplifiers: tools for discarding unimportant aspects of a problem and for selecting the important aspects.

These reasoning tools are of three types:

1. **Organizing complexity**
   - Divide and conquer
   - Abstraction
2. **Lossless compression**
   - Symmetry and conservation
   - Proportional reasoning
   - Dimensions
3. **Lossy compression**

- Easy cases
- Probabilistic reasoning
- Lumping
- Spring models

The first type of tool helps manage complexity. The second type helps remove complexity that is merely apparent. The third type helps discard complexity.

With these tools we explore the natural and manmade worlds, using examples from diverse fields such as quantum mechanics, general relativity, mechanical engineering, biophysics, recreational mathematics, and climate change. This diversity has two purposes. First, the diversity shows how a small toolbox can explain important features of the manmade and engineered worlds. The diversity provides a library of models for your own analyses.

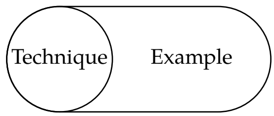

Second, the diversity separates the tool from the details of its use. A tool is difficult to appreciate abstractly, without an example. However, if you see only one use of a tool, the tool is difficult to distinguish from the example. An expert, familiar with the tool, knows where the idea ends and the details begin. But when you first learn a tool, you need to learn the boundary.

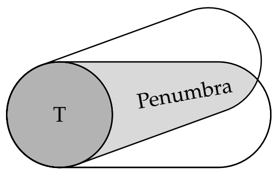

An answer is a second example. To the extent that the second example is similar to the first, the tool plus first use overlaps the tool plus second use. The overlap includes a penumbra around the tool. The penumbra is smaller than it is with only one example: Two uses delimit the boundaries of the tool more clearly than one example does.

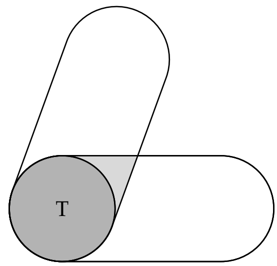

More clarity comes using an example from a distant field. The penumbra shrinks, which separates the tool from examples of its use. For example, using dimensional analysis in a physics problem and an economics analysis clarifies what part of the illustration is specific to physics or economics and what part is transferable to other problems. Focus on the transferable ideas; they are useful in any career!

This book is designed for self study. Therefore, please try the problems. The problems are of two types. The first type are problems marked with a wedge in the margin. They are breathers during an analysis: a place to develop your understanding by working out the next steps in an analysis. Those problems are answered in the subsequent text where you can check your thinking and my analysis – please let me know of any errors! The second type of problem, the numbered problems, give practice with the tools, extend a derivation, or develop a useful or enjoyable model. Most numbered problems have answers at the end of the book.

I hope that you find the tools, problems, and models useful in your career. And I hope that the diversity of examples connects with and aids your curiosity about how the world is put together.

Bon voyage!

Part 1

Organizing complexity

The first solution to the messiness and complexity of the world, just as with the mess on our desks and in our living spaces, is to organize the complexity. Two techniques for organizing complexity are the subject of Part 1.

The first technique is divide-and-conquer reasoning: dividing a large problem into manageable subproblems. The second technique is abstraction: choosing compact representations that hide unimportant details in order to reveal important features. The next two chapters illustrate these techniques with many examples.

1
Divide and conquer

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [16, Book VI]:

> A Captain ought…endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces,* and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have $2,789,164$ files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general, divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

▶ *How far apart are the pits on a compact disc (CD) or CDROM?*

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

That calculation was simplified by rounding up the number of bits from $5{\cdot}10^9$ to 10^{10}. The factor of 2 increase means that 1 μm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 μm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.

2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.

3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4\times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4\cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4\cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?

In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage

In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(\text{few})^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [8].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

   Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from $0.35\,\mu\text{m}$ to $0.7\,\mu\text{m}$ – let's call it $0.5\,\mu\text{m}$ – I estimate the pit spacing to be $1\,\mu\text{m}$.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – $1\,\mu\text{m}$. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: $1.4\,\mu\text{m}$ (capacity), $1\,\mu\text{m}$ (optics), and $1\,\mu\text{m}$ (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is $0.1\,\mu\text{m}$. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition

The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy

Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

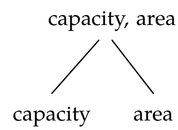

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

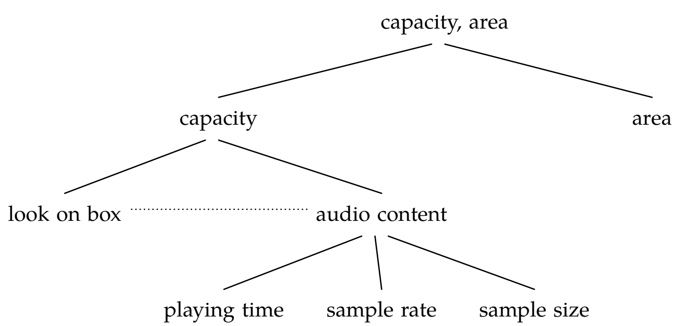

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\text{cm})^2$.

Here is the quantified tree:

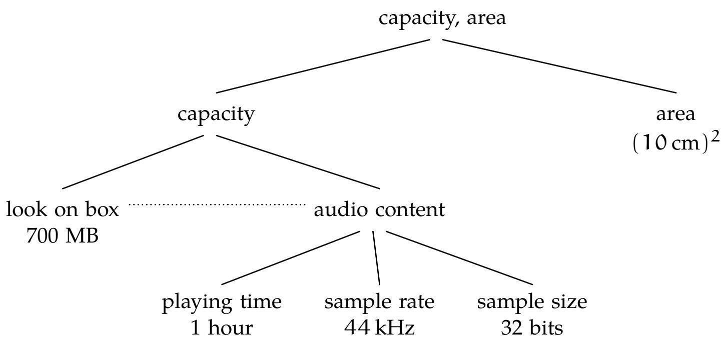

The final step is to propagate estimates upward, from children to parent, until reaching the root.

Draw the resulting tree.

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly $1.4\,\mu$m.

Propagating estimates from leaf to root gives the following tree:

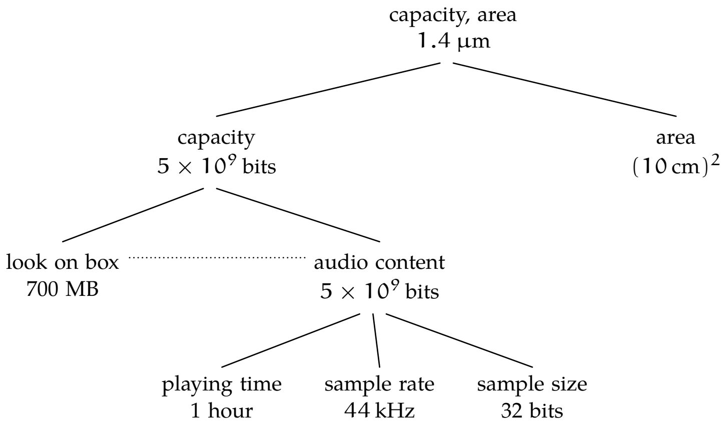

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

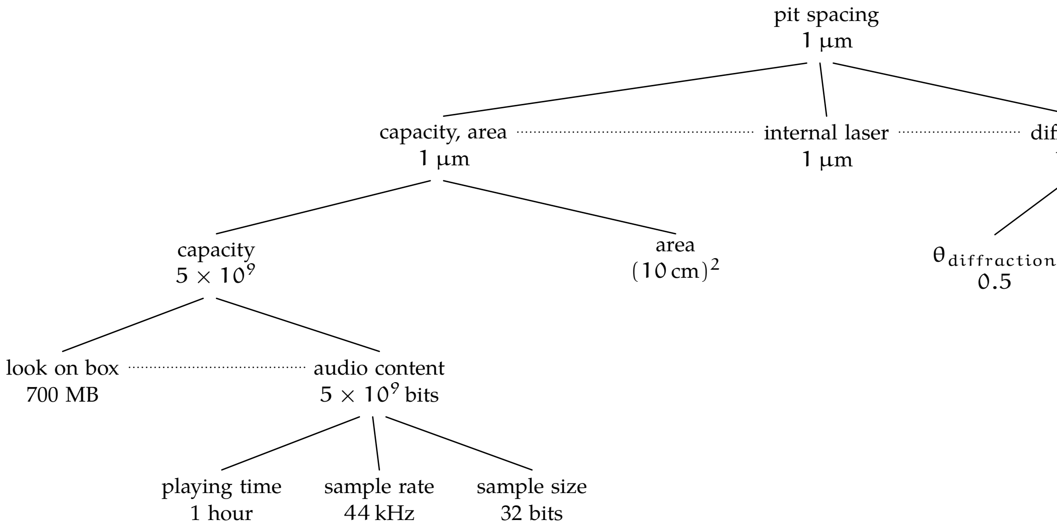

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

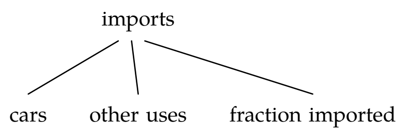

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

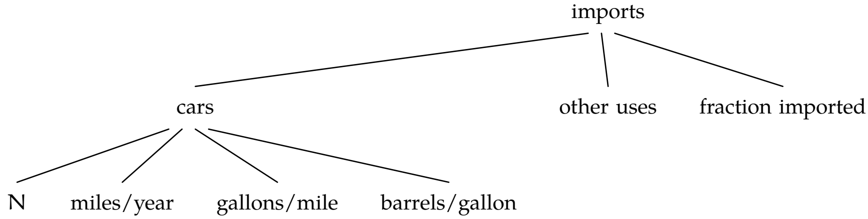

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about \$100, and a gallon of gasoline costs about \$2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

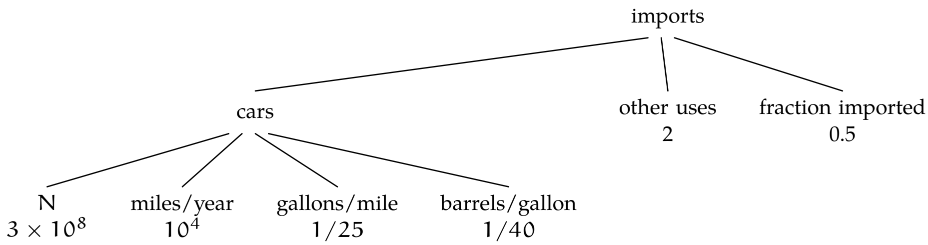

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year.}$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year.}$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25% higher than the estimate!

Problem 1.5 Midpoints

The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [9], a valuable book for anyone interested in how to design large systems. The author isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 8.2). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing

documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.

9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

▶ *How can you use these building blocks to print the 23rd line of a file?*

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23$^{\text{rd}}$ line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

   ```
   head -23 | tail -1
   ```

   therefore reads lines typed at the keyboard, prints the 23rd line, and exits (even if the user is still typing).

2. Tell `head` to read its input from a file – for example from an English dictionary. On my GNU/Linux computer, the English dictionary is the file `/usr/share/dict/words`. It contains one word per line, so the following pipeline prints the 23rd word from the dictionary:

   ```
   head -23 /usr/share/dict/words | tail -1
   ```

3. Let `head` read from its standard input, but connect the standard input to a file:

   ```
   head -23 < /usr/share/dict/words | tail -1
   ```

   The `<` operator tells the UNIX command interpreter to connect the file `/usr/share/dict/words` to the input of `head`. The system tricks `head` into thinking its reading from the keyboard, but the input comes from the file – without requiring any change in the program!

4. Use the `cat` program to achieve the same effect as the preceding method. The `cat` program copies its input file(s) to the output. This extended pipeline therefore has the same effect as the preceding method:

   ```
   cat /usr/share/dict/words | head -23 | tail -1
   ```

   This longer pipeline is slightly less efficient than using the redirection operator. The pipeline requires an extra program (`cat`) copying its input to its output, whereas the redirection operator lets the lower level of the UNIX system achieve the same effect (replumbing the input) without the gratuitous copy.

As practice, let's use the UNIX approach to divide and conquer a search problem:

Imagine a dictionary of English alphabetized from right to left instead of the usual left to right. In other words, the dictionary begins with words that end in 'a'. In that dictionary, what word immediately follows `trivia`?

This whimsical problem is drawn from a scavenger hunt [24]created by the computer scientist Donald Knuth, whose many accomplishments include the TEX typesetting system used to produce this book.

The UNIX approach divides the problem into two parts:

1. Make a dictionary alphabetized from right to left.
2. Print the line following 'trivia'.

The first problem subdivides into three parts:

1. Reverse each line of a regular dictionary.
2. Alphabetize (sort) the reversed dictionary.
3. Reverse each line to undo the effect of step 1.

The second part is solved by the UNIX utility `sort`. For the first and third parts, perhaps a solution is provided by an item in UNIX toolbox. However, it would take a long time to thumb through the toolbox hoping to get lucky: My computer tells me that it has over 8000 system programs.

Fortunately, the UNIX utility `man` does the work for us. `man` with the `-k` option, with the 'k' standing for keyword, lists programs with a specified keyword in their name or one-line description. On my laptop, `man -k reverse` says:

```
  $ man -k reverse
  col (1)              - filter reverse line feeds from in-
put
  git-rev-list (1)     - Lists commit objects in reverse chrono-
logical order
  rev (1)              - reverse lines of a file or files
  tac (1)              - concatenate and print files in re-
verse
  xxd (1)              - make a hexdump or do the reverse.
```

Understanding the free-form English text in the one-line descriptions is not a strength of current computers, so I leaf through this list by hand – but it contains only five items rather than 8000. Looking at the list, I spot `rev` as a filter that reverses each line of its input.

How do you use `rev` and `sort` to alphabetize the dictionary from right to left?

Therefore the following pipeline alphabetizes the dictionary from right to left:

```
rev < /usr/share/dict/words | sort | rev
```

The second problem – finding the line after 'trivia' – is a task for the pattern-searching utility `grep`. If you had not known about `grep`, you might find it by asking the system for help with `man -k pattern`. Among the short list is

```
grep (1)             - print lines matching a pattern
```

In its simplest usage, `grep` prints every input line that matches a specified pattern. For example,

```
grep 'trivia' < /usr/share/dict/words
```

prints all lines that contain `trivia`. Besides `trivia` itself, the output includes `trivial`, `nontrivial`, `trivializes`, and similar words. To require that the word match `trivia` with no characters before or after it, give grep this pattern:

```
grep '^trivia$' < /usr/share/dict/words
```

The patterns are regular expressions. Their syntax can become arcane but their important features are simple. The ˆ character matches the beginning of the line, and the $ character matches the end of the line. So the pattern t̂rivia$ selects only lines that contain exactly the text `trivia`.

This invocation of `grep`, with the special characters anchoring the beginning and ending of the lines, simply prints the word that I specified. How could such an invocation be useful?

That invocation of `grep` tells us only that `trivia` is in the dictionary. So it is useful for checking spelling – the solution to a problem, but not to our problem of finding the word that follows `trivia`. However, Invoked with the `-A` option, `grep` prints lines following each matching line. For example,

```
grep -A 3 '^trivia$' < /usr/share/dict/words
```

will print 'trivia' and the three lines (words) that follow it.

```
trivia
trivial
trivialities
triviality
```

To print only the word after 'trivia' but not 'trivia' itself, use `tail`:

```
grep -A 1 '^trivia$' < /usr/share/dict/words | tail -1
```

These small solutions combine to solve the scavenger-hunt problem:

```
rev </usr/share/dict/words | sort | rev | grep -A 1 '^trivia$'
| tail -1
```

▶ *Try it on a local UNIX or GNU/Linux system. How well does it work?*

Alas, on my system, the pipeline fails with the error

```
  rev: stdin: Invalid or incomplete multibyte or wide char-
acter
```

The `rev` program is complaining that it does not understand a character in the dictionary. `rev` is from the old, ASCII-only days of UNIX, when each character was limited to one byte; the dictionary, however, is a modern one and includes Unicode characters to represent the accented letters prevalent in European languages.

To solve this unexpected problem, I clean the dictionary before passing it to `rev`. The cleaning program is again the filter `grep` told to allow through only pure ASCII lines. The following command filters the dictionary to contain words made only of unaccented, lowercase letters.

```
grep '^[a-z]*$' < /usr/share/dict/words
```

This pattern uses the most important features of the regular-expression language. The ^ and $ characters have been explained in the preceding examples. The `[a-z]` notation means 'match any character in the range a to z – i.e. match any lowercase letter.' The `*` character means 'match zero or more occurrences of the preceding regular expression'. So ^ [a-z]*$ matches any line that contains only lowercase letters – no Unicode characters allowed.

The full pipeline is

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

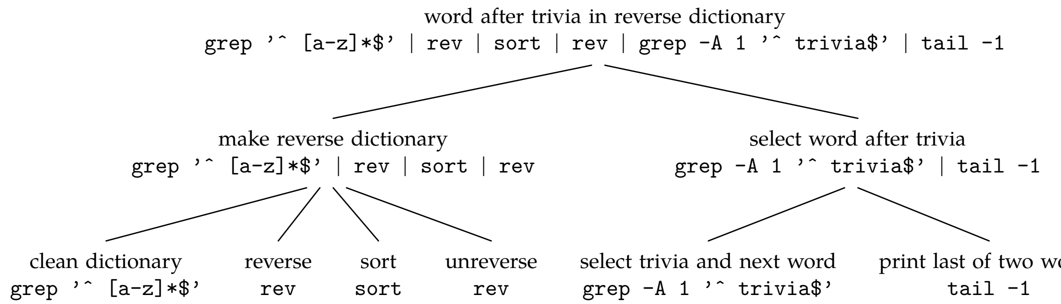

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

By breaking hard problems into comprehensible units, the divide-and-conquer tool helps us organize complexity. The next chapter examines its cousin abstraction, another way to organize complexity.

Problem 1.7 Air mass
Estimate the mass of air in the 6.055J/2.038J classroom and explain your estimate with a tree. If you have not seen the classroom yet, then make more effort to come to lecture (!); meanwhile pictures of the classroom are linked from the course website.

Problem 1.8 747
Estimate the mass of a full 747 jumbo jet, explaining your estimate using a tree. Then compare with data online. We'll use this value later this semester for estimating the energy costs of flying.

Problem 1.9 Random walks and accuracy of divide and conquer
Use a coin, a random-number function (in whatever programming language you like), or a table of reasonably random numbers to do the following experiments or their equivalent.

The first experiment:

1. Flip a coin 25 times. For each heads move right one step; for each tails, move left one step. At the end of the 25 steps, record your position as a number between -25 and 25.
2. Repeat the above procedure four times (i.e. three more times), and mark your four ending positions on a number line.

The second experiment:

1. Flip a coin once. For heads, move right 25 steps; for tails, move left 25 steps.
2. Repeat the above procedure four times (i.e. three more times), and mark your four ending positions on a second number line.

Compare the marks on the two number lines, and explain the relation between this data and the model from lecture for why divide and conquer often reduces errors.

Problem 1.10 Fish tank

Estimate the mass of a typical home fish tank (filled with water and fish): a useful exercise before you help a friend move who has a fish tank.

Problem 1.11 Bandwidth

Estimate the bandwidth (bits/s) of a 747 crossing the Atlantic filled with CDROM's.

Problem 1.12 Repainting MIT

Estimate the cost to repaint all indoor walls in the main MIT classroom buildings. [with thanks to D. Zurovcik]

Problem 1.13 Explain a UNIX pipeline

What does this pipeline do?

```
ls -t | head | tac
```

[Hint: If you are not familiar with UNIX commands, use the `man` command on any handy UNIX or GNU/Linux system.]

Problem 1.14 Design a UNIX pipeline

Make a pipeline that prints the ten most common words in the input stream, along with how many times each word occurs. They should be printed in order from the the most frequent to the less frequent words. [Hint: First **tr**anslate any non-alphabetic character into a newline. Useful utilities include `tr` and `uniq`.]

2
Abstraction

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

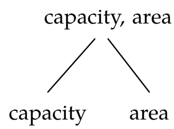

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.3.1:

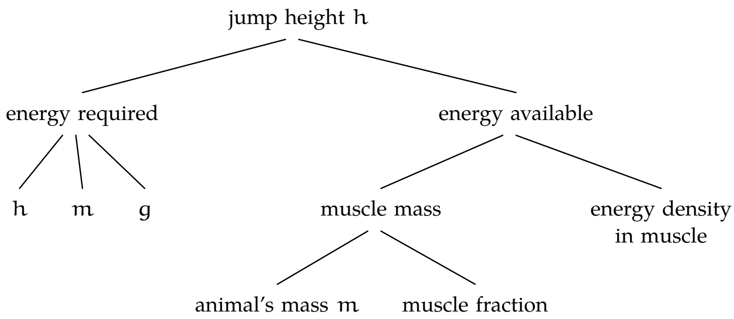

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams as abstractions

Diagrams are themselves a powerful kind of abstraction. Diagrams are an abstraction because they force one to discard irrelevant details, reducing a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation, the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

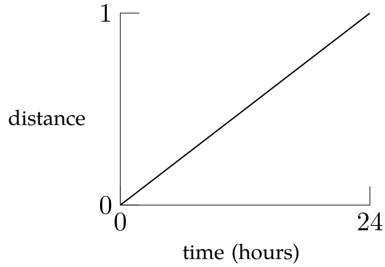

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

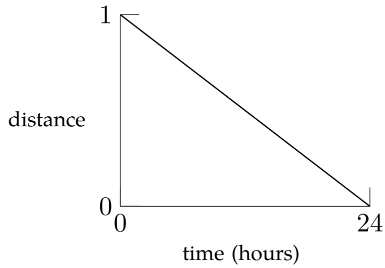

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

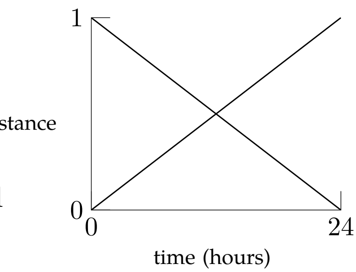

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

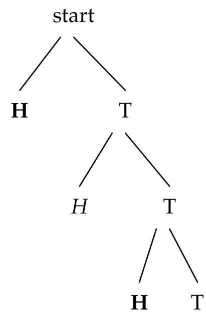

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and $p/2$ add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \tag{2.2}$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \tag{2.3}$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \tag{2.4}$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \tag{2.5}$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \tag{2.6}$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \; \Big| \; 3 \times 7 + 5 \times 2 \; \Big| \; 5 \times 7. \tag{2.7}$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where $y|x$ means $100y + x$, the product 3247×1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \;\Big|\; 32 \times 98 + 47 \times 17 \;\Big|\; 47 \times 98. \tag{2.8}$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal $y|x = 10y + x$ notation),

$$3|2 \times 1|7 = 3 \times 2 \;\Big|\; 3 \times 7 + 2 \times 1 \;\Big|\; 2 \times 7. \tag{2.9}$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \;\Big|\; (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \;\Big|\; a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35×27 example.

$$3|5 \times 2|7 = 3 \times 2 \;\Big|\; (3 + 5)(2 + 7) - 3 \times 2 - 5 \times 7 \;\Big|\; 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6\,|\,31\,|\,35 = 6 \times 100 + 31 \times 10 + 35 = 945, \tag{2.10}$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [13]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log \log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm

Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

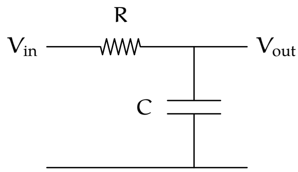

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an *RC* circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the *RC* system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geq 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

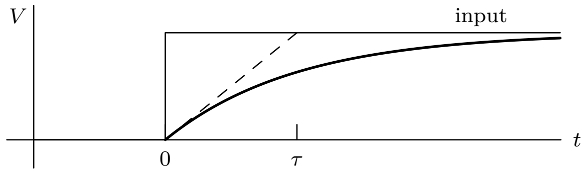

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by 90° relative to the input oscillations (where 360° is a full period). Why 90°? In the complex plane, dividing by i is equivalent to rotating clockwise by 90°. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

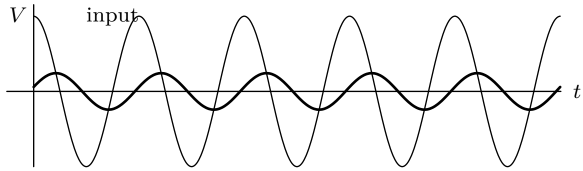

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\ \tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.5\,\text{day}}_{\tau}, \qquad (2.18)$$

or approximately 3. In the (South African) winter, the outside temperature varied between 45 °F and 75 °F. This 30 °F outside variation gets shrunk by a factor of 3, giving an inside variation of 10 °F. This variation occurred around the average outside temperature of 60 °F, so the inside temperature varied between 55 °F and 65 °F. Furthermore, if the coldest outside temperature is at midnight, the coldest inside temperature is delayed by almost 6 hr (the one-quarter-period delay). Indeed, the house did feel coldest early in the morning, just as I was getting up – as predicted by this simple model of heat flow that is based on a circuit-analysis abstraction.

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in a famous representation: Richard Feynman's diagrams for calculations in the theory of quantum electrodynamics (how radiation interacts with matter). Those diagrams are discussed in [32]and [12].

The main ideas in this chapter:

For more on the value of diagrams, see [28]and [19].

Problem 2.3 Spacetime diagrams

Learn about spacetime diagrams. My favorite source is *Spacetime Physics* [30].

Problem 2.4 Word processors

Compare WYSIWIG (what you see is what you get) word processors such as WordPerfect or Microsoft Word with document formatting systems such as TEX or ConTEXt (used to typeset this book).

Problem 2.5 Longest left-handed word

What is the longest word in the dictionary that can be typed with only the left hand (on a qwerty keyboard)?

Part 2

Lossless compression

The first part discussed methods for organizing and therefore for managing complexity. The remaining two parts discuss how to *discard complexity*. The discarded complexity can be actual complexity (Part 3) or it can be only apparent complexity – whereupon discarding it does not discard information. Such lossless compression is the subject of this part.

The three methods are symmetry and conservation, proportional reasoning, and dimensional analysis. Proportional reasoning and dimensional analysis are, additionally, examples of symmetry reasoning. Therefore, the next chapter introduces symmetry and conservation reasoning.

3
Symmetry and conservation

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

In this form, $2S$ is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

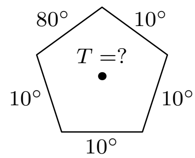

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa \nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa \nabla^2 T = 0$.

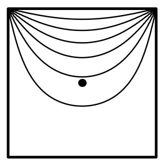

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at $10\,^\circ$, $10\,^\circ$, $10\,^\circ$, and $80\,^\circ$ (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every $10\,^\circ$). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

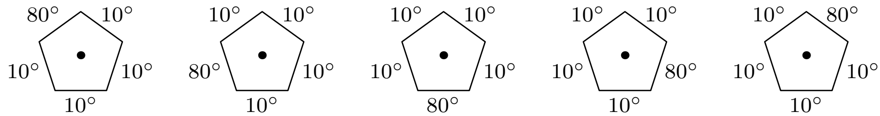

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

3.2 Cube solitaire

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

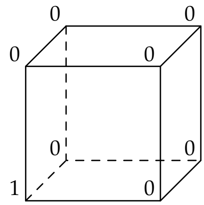

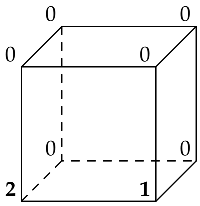

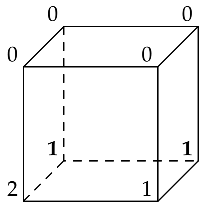

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

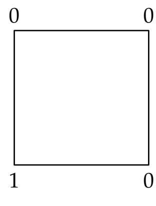

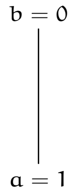

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

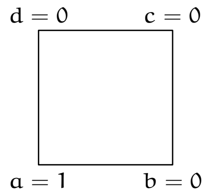

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

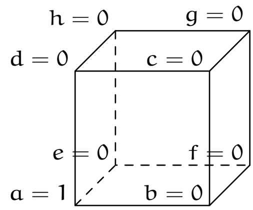

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a-b+c-d+f-g+h-e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

$$a-b+c-d+f-g+h-e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

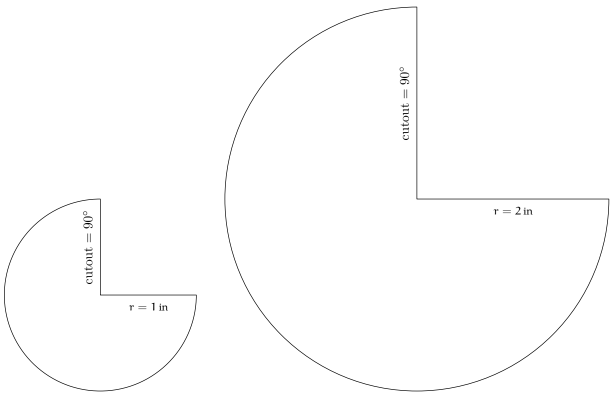

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds $v_{\text{big}}/v_{\text{small}}$? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [7, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

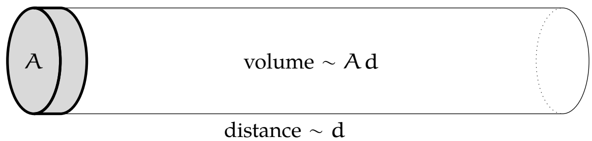

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is ρAd, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is ρAv^2d, making the drag force

$$F \sim \rho Av^2.$$

The analysis has a divide-and-conquer tree:

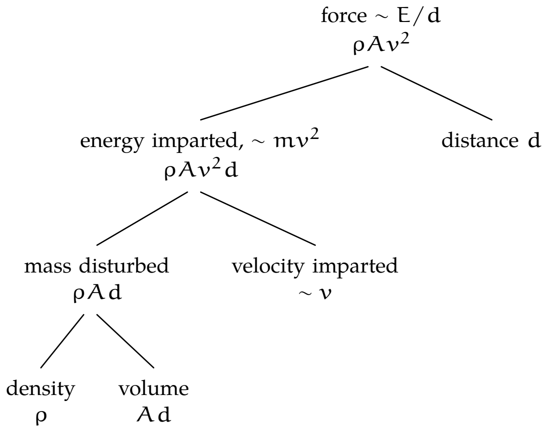

The result that $F_{\text{drag}} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 7.10), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{\text{drag}}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{\text{drag}} = F_{\text{drag}} v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{\text{drag}} = P_{\text{athlete}}$ gives

$$v_{\text{max}} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4\,\text{m}$ (a typical house is $3\,\text{m}$ per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5\,\text{s}$ for the climb. My mass is $60\,\text{kg}$, so my power output was

$$P_{\text{author}} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

P_{athlete} should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{\text{athlete}} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

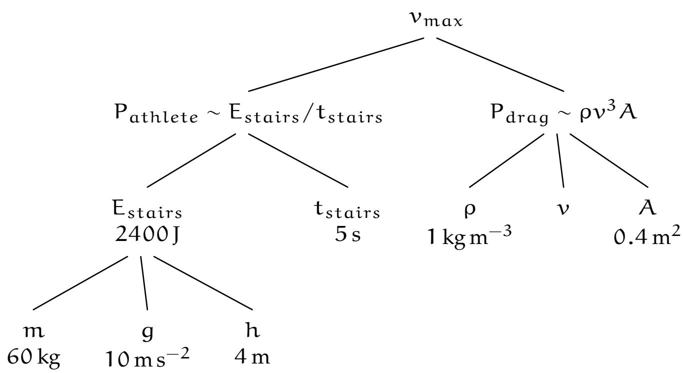

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{\text{max}} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{\text{max}} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

So the world record should be, if this analysis has any correct physics in it, around $10\,\mathrm{m\,s^{-1}}$ or 22 mph.

The world one-hour record – where the contestant cycles as far as possible in one hour – is 49.7 km or 30.9 mi. The estimate based on drag is reasonable!

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{\text{air}} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

$$m_{\text{air}} v_{\text{down}} \sim \frac{Mgs}{v}$$

so

$$v_{\text{down}} \sim \frac{Mgs}{v m_{\text{air}}} \sim \frac{Mgs}{\rho v L^2 s} = \frac{Mg}{\rho v L^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{\text{air}} v_{\text{down}}^2$. That energy factors into $(m_{\text{air}} v_{\text{down}}) v_{\text{down}}$, so

$$E_{\text{lift}} \sim \underbrace{m_{\text{air}} v_{\text{down}}}_{Mgs/v} v_{\text{down}} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho v L^2}}_{v_{\text{down}}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{\text{drag}} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

$$F_{\text{drag}} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{\text{drag}} = C \rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{\text{drag}} = C \rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{\text{lift}}} + \underbrace{C \rho v^2 L^2 s}_{E_{\text{drag}}}.$$

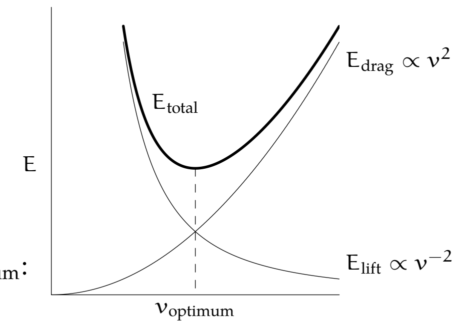

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed v_{optimum}: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2} \rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{\text{drag}} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{\text{drag}} = \frac{1}{2} C' A_{\text{wing}} \rho v^2.$$

Alas, this formula is a third convention for drag coefficients, depending on whether the drag is referenced to the cross-sectional area A, wing area A_{wing}, or squared wingspan L^2.

It is easy to convert between the definitions. Just equate the standard definition

$$F_{\text{drag}} = \frac{1}{2} C' A_{\text{wing}} \rho v^2.$$

to our definition

$$F_{\text{drag}} = C L^2 \rho v^2$$

to get

$$C = \frac{1}{2} \frac{A_{\text{wing}}}{L^2} C' = \frac{1}{2} \frac{l}{L} C',$$

since $A_{\text{wing}} = Ll$ where l is the wing width. For a 747, $l \sim 10\,\text{m}$ and $L \sim 60\,\text{m}$, so $C \sim 1/600$.

Combine the values to find the range:

$$s \sim \frac{\beta \mathcal{E}}{C^{1/2} g} \sim \frac{0.4 \times 10^7\,\text{J}\,\text{kg}^{-1}}{(1/600)^{1/2} \times 10\,\text{m}\,\text{s}^{-2}} \sim 10^7\,\text{m} = 10^4\,\text{km}.$$

The maximum range of a 747-400 is 13,450 km, so the approximate analysis of the range is unreasonably accurate.

Problem 3.1 Integrals

Evaluate these definite integrals:

a. $\int_{-10}^{10} x^3 e^{-x^2}\, dx$

b. $\int_{-\infty}^{\infty} \frac{x^3}{1 + 7x^2 + 18x^8}\, dx$

Problem 3.2 Number sum

Use symmetry to find the sum of the integers between 200 and 300 (inclusive).

Problem 3.3 Heat equation

In lecture we used symmetry to argue that the temperature at the center of the metal sheet is the average of the temperatures of the sides.

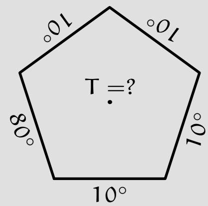

Check this result by making a simulation or, if you are bold but crazy, by finding an analytic solution of the heat equation.

Problem 3.4 Symmetry for algebra

Use symmetry to find $(a - b)^3$.

Problem 3.5 Symmetry for second-order systems

This problem analyzes the frequency of maximum gain for an *LRC* circuit or, equivalently, for a damped spring–mass system. The gain of such a system is the ratio of the input amplitude to the output amplitude as a function of frequency.

If the output voltage is measured across the resistor, and you drive the circuit with a voltage oscillating at frequency ω, the gain is (in a suitable system of units):

$$G(\omega) = \frac{j\omega}{1 + j\omega/Q - \omega^2},$$

where $j = \sqrt{-1}$ and Q is quality factor, a dimensionless measure of the damping.

Do not worry if you do not know where that gain formula comes from. The purpose of this problem is not its origin, but rather using symmetry to maximize its magnitude.

a. Show that the magnitude of the gain is

$$|G(\omega)| = \frac{\omega}{\sqrt{\left(1 - \omega^2\right)^2 + \omega^2/Q^2}}.$$

b. Find a variable substitution (a symmetry operation) $\omega_{\text{new}} = f(\omega)$ that turns $|G(\omega)|$ into $|H(\omega_{\text{new}})|$ such that G and H are the same function (i.e. they have the same structure but with ω in G replaced by ω_{new} in H).

c. Use the form of that symmetry operation to maximize $|G(\omega)|$ without using calculus.

d. *[Optional, for masochists!]* Maximize $|(G\omega)|$ using calculus.

Problem 3.6 Inertia tensor

[For those who know about inertia tensors.] Here is the inertia tensor (the generalization of moment of inertia) of a particular object, calculated in a lousy coordinate system:

$$\begin{pmatrix} 4 & 0 & 0 \\ 0 & 5 & 4 \\ 0 & 4 & 5 \end{pmatrix}$$

a. Change coordinate systems to a set of principal axes. In other words, write the inertia tensor as

$$\begin{pmatrix} I_{\text{xx}} & 0 & 0 \\ 0 & I_{\text{yy}} & 0 \\ 0 & 0 & I_{\text{zz}} \end{pmatrix}$$

and give I_{xx}, I_{yy}, and I_{zz}. *Hint:* What properties of a matrix are invariant when changing coordinate systems?

b. Give an example of an object with a similar inertia tensor. On Friday in class we'll have a demonstration.

Problem 3.7 Resistive grid

In an infinite grid of 1-ohm resistors, what is the resistance measured across one resistor?

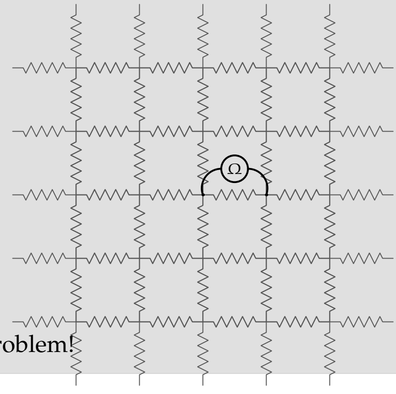

To measure resistance, an ohmmeter injects a current I at one terminal (for simplicity, say $I = 1\,\text{Å}$), removes the same current from the other terminal, and measures the resulting voltage difference V between the terminals. The resistance is $R = V/I$.

Hint: Use symmetry. But it's still a hard problem!

4
Proportional reasoning

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having fed all summer is perhaps 30 or 40% fat. So $\beta_{\text{plane}}/\beta_{\text{bird}} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{\text{plane}}/\mathcal{E}_{\text{bird}} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{\text{bird}}/C_{\text{plane}} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is $11,570\,\text{km}$, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is $13,450\,\text{km}$, only slightly longer than the godwit's range.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

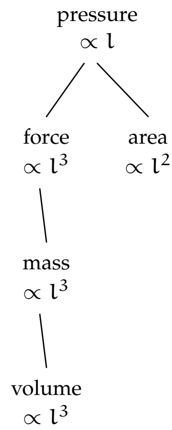

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

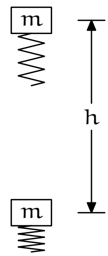

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{\text{jump}} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{\text{muscle}} \sim \alpha m$$

or, as a proportionality,

$$m_{\text{muscle}} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{\text{avail}} \sim \mathcal{E} m_{\text{muscle}}$$

or, as a proportionality,

$$E_{\text{avail}} \propto m_{\text{muscle}},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

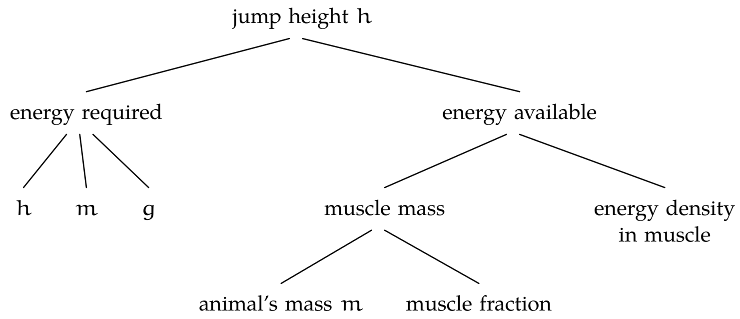

Now finish propagating toward the root. The available energy is

$$E_{\text{avail}} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{\text{avail}} \propto m$$

and the required energy is

$$E_{\text{required}} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

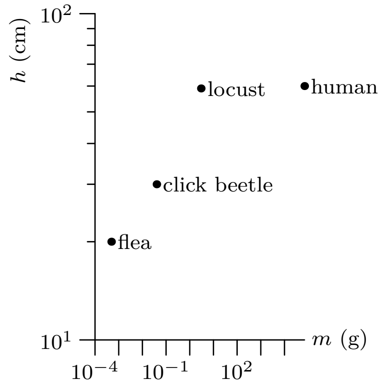

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^4$	60

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

$$E_{\text{drag}} \sim \underbrace{\rho v^2 A}_{F_{\text{drag}}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{\text{drag}}}{E_{\text{required}}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and ρAh is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator ρAh is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{\text{drag}}}{E_{\text{required}}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{\text{drag}}}{E_{\text{required}}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{\text{animal}} l^3}.$$

It simplifies to

$$\frac{E_{\text{drag}}}{E_{\text{required}}} \sim \frac{\rho}{\rho_{\text{animal}}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\mathrm{kg\,m^{-3}}$. The density of an animal is roughly the density of water, so $\rho_{\mathrm{animal}} \sim 10^3\,\mathrm{kg\,m^{-3}}$. The typical jump height – which is where the data substitutes for the constant of proportionality – is 60 cm or roughly 1 m. A flea's length is about 1 mm or $l \sim 10^{-3}\,\mathrm{m}$. So

$$\frac{E_{\mathrm{drag}}}{E_{\mathrm{required}}} \sim \frac{1\,\mathrm{kg\,m^{-3}}}{10^3\,\mathrm{kg\,m^{-3}}}\,\frac{1\,\mathrm{m}}{10^{-3}\,\mathrm{m}} \sim 1.$$

The ratio being unity means that if a flea would jump to 60 cm, overcoming drag would require roughly as much as energy as would the jump itself in vacuum.

Drag provides a plausible explanation for why fleas do not jump as high as the typical height to which larger animals jump.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{\text{max}} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{\text{swimmer}}/v_{\text{cyclist}}$, and then use that ratio along with v_{cyclist} to work out v_{swimmer}.

The speed v_{max} is

$$v_{\text{max}} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{\text{swimmer}}}{v_{\text{cyclist}}} \sim \left(\frac{P_{\text{swimmer}}}{P_{\text{cyclist}}}\right)^{1/3} \times \left(\frac{\rho_{\text{swimmer}}}{\rho_{\text{cyclist}}}\right)^{-1/3} \times \left(\frac{A_{\text{swimmer}}}{A_{\text{cyclist}}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\text{m}\,\text{s}^{-1}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\text{m}\,\text{s}^{-1}$, significantly higher than the prediction of $1\,\text{m}\,\text{s}^{-1}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600\,\text{mph}$ and $v_{\text{car}} \sim 60\,\text{mph}$, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

$$A_{\text{plane}} \sim 6\,\text{m} \times 6\,\text{m} = 36\,\text{m}^2,$$

whereas

$$A_{\text{car}} \sim 2\,\text{m} \times 1.5\,\text{m} = 3\,\text{m}^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{\text{drag}} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

More on proportional reasoning

Problem 4.1 Raindrop speed

a. How does a raindrop's terminal velocity v depend on the raindrop's radius r?

b. Estimate the terminal speed for a typical raindrop.

c. How could you check your estimate in part (b)?

Problem 4.2 Mountains

Look up the height of the tallest mountain on earth, Mars, and Venus, and explain any pattern in the three values.

Problem 4.3 Highway vs city driving

In lecture we derived a measure of how important drag is for a car moving at speed v for a distance d:

$$\frac{E_{\text{drag}}}{E_{\text{kinetic}}} \sim \frac{\rho v^2 A d}{m_{\text{car}} v^2}.$$

a. Show that the ratio is equivalent to the ratio

$$\frac{\text{mass of the air displaced}}{\text{mass of the car}}$$

and to the ratio

$$\frac{\rho_{\text{air}}}{\rho_{\text{car}}} \times \frac{d}{l_{\text{car}}},$$

where ρ_{car} is the density of the car (i.e. its mass divided by its volume) and l_{car} is the length of the car.

b. Make estimates for a typical car and find the distance d at which the ratio becomes significant (say, roughly 1). How does the distance compare with the distance between exits on the highway and between stop signs or stoplights on city streets?

Problem 4.4 Gravity on the moon

In this problem you use a scaling argument to estimate the strength of gravity on the surface of the moon.

a. Assume that a planet is a uniform sphere. What is the proportionality between the gravitational acceleration g at the surface of a planet and the planet's radius R and density ρ?

b. Write the ratio $g_{\text{moon}}/g_{\text{earth}}$ as a product of dimensionless factors as in the analysis of the fuel efficiency of planes.

c. Estimate those factors and estimate the ratio $g_{\text{moon}}/g_{\text{earth}}$, then estimate g_{moon}. [Hint: To estimate the radius of the moon, whose angular size you can estimate by looking at it, you might find it useful to know that the moon is $4 \cdot 10^8$ m distant from the earth.]

d. Look up g_{moon} and compare the value to your estimate, venturing an explanation for any discrepancy.

Problem 4.5 Checking plane fuel-efficiency calculation

This problem offers two more methods to estimate the fuel efficiency of a plane.

a. Use the cost of a plane ticket to estimate the fuel efficiency of a 747, in passenger–miles per gallon.

b. According to Wikipedia, a 747-400 can hold up to $2{\cdot}10^5\,\ell$ of fuel for a maximum range of $1.3 \cdot 10^4$ km. Use that information to estimate the fuel efficiency of the 747, in passenger–miles per gallon.

How do these values compare with the rough result from lecture, that the fuel efficiency is comparable to the fuel efficiency of a car?

5
Dimensions

5.1 Power of multinational corporations

Critics of globalization often make the following comparison [17] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains steady from year to year) would be roughly \$1 trillion per decade and be reported as \$1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes \$2 billion per week, reported as \$2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [18, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any meanginful statement about these variables looks like

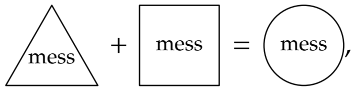

where the various messes mean 'a horrible combination of E, s, v, ρ, L, and m.

As horrible as that statement is, it permits the following rewriting: Divide each term by the first one (the triangle). Then

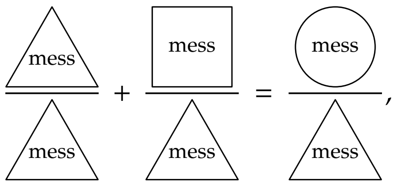

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

So the rewritten expression is dimensionless. Nothing in the rewriting depended on the particular form of the statement, except that each term has the same dimensions.

Therefore, *any meaningful statement can be rewritten in dimensionless form.*

Dimensionless forms are made from dimensionless ratios, so all you need are dimensionless ratios, and you can do all your thinking with them. As a negative example, revisit the comparison between Exxon's net worth and Nigeria's GDP (Section 5.1). The dimensions of net worth are simply money. The dimensions of GDP are money per time. These two quantities cannot form a dimensionless group! With just these two quantities, no meanginful statements are possible.

Here is a further example to show how this change simplifies your thinking. This example uses familiar physics so that you can concentrate on the new idea of dimensionless ratios.

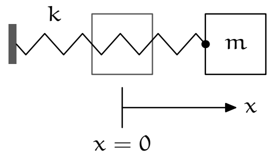

The problem is to find the period of an oscillating spring–mass system given an initial displacement x_0, then allowed to oscillate freely. The relevant variables that determine the period T are mass m, spring constant k, and amplitude x_0. Those three variables completely describe the system, so any true statement about period needs only those variables.

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = \mathrm{MLT}^{-2}$ and $[x] = \mathrm{L}$,

$$[k] = \mathrm{MT}^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, , \ldots.$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.
2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [20, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [7, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [25].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = \left[r^2\right] \times [F] = \mathrm{ML^3T^{-2}}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = \mathrm{ML^2T^{-1}}.$$

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	$\mathrm{ML^3T^{-2}}$	
m_e	M	electron mass
$\hbar$	$\mathrm{ML^2T^{-1}}$	quantum

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = \mathrm{ML},$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant.}$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\mathrm{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

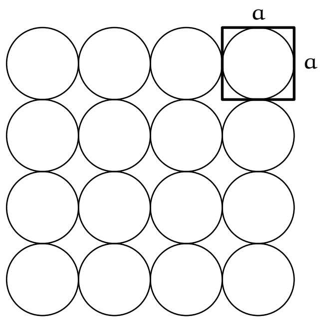

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of $N_{\text{A}}/N_{\text{A}}$, where N_{A} is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_{\text{A}}}{(3\,\text{Å})^3 \times N_{\text{A}}}.$$

The numerator is A g, because that is how N_{A} is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_{A}. However, N_{A} is more familiar than m_{proton} because N_{A} arises in chemistry and physics. Using N_{A} also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

The table compares the estimate against reality. Most everyday elements have atomic masses between 15 and 150, so the density estimate explains why most densities lie between 1 and 10 g cm^{-3}. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is 3 g cm^{-3}, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid 3000 kg m^{-3}. Numbers such as 3000 are unwieldy. Each time you use it, you have to think, 'How many powers of ten were there again?' So the table tabulates densities using the cgs units of g cm^{-3}. I even threw a joker into the pack – water is not an element! – but the density estimate is amazingly accurate.

Element	$\rho_{\text{estimated}}$	ρ_{actual}
Li	0.39	0.54
H_2O	1.0	1.0
Si	1.56	2.4
Fe	3.11	7.9
Hg	11.2	13.5
Au	10.9	19.3
U	13.3	18.7

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

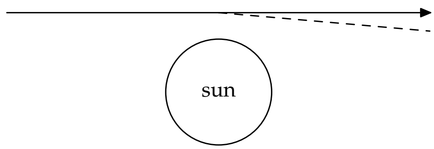

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

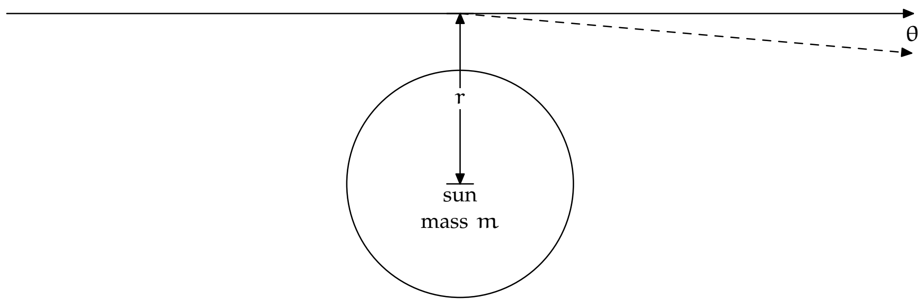

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$\mathrm{L^3T^{-2}M^{-1}}$
r	distance from center of sun	L
c	speed of light	$\mathrm{LT^{-1}}$

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

$$\theta = 0.3\frac{Gm}{rc^2}$$

are also possible. This freedom means that we should use a twiddle rather than an equals sign:

$$\theta \sim \frac{Gm}{rc^2}.$$

5.4.4 Comparison with exact calculations

All theories of gravity have the same form for the result, namely

$$\theta \sim \frac{Gm}{rc^2}.$$

The difference among the theories is in the value for the missing dimensionless constant:

$$\theta = \frac{Gm}{rc^2} \times \begin{cases} 1 & \text{(simplest guess);} \\ 2 & \text{(Newtonian gravity);} \\ 4 & \text{(Einstein's theory).} \end{cases}$$

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just the simplest possible guess. The 2 for Newtonian gravity is from integrating angular factors like cosine and sine that determine the position of the photon as it moves toward and past the sun.

The most interesting constant is the 4 for general relativity, which is double the Newtonian value. The fundamental reason for the factor of 2 is that special relativity puts space and time on an equal footing to make spacetime. The theory of general relativity builds on special relativity by formulating gravity as curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; but general relativity also handles the space curvature. Most objects move much slower than the speed of light, so they move much farther in time than in space and see mostly the time curvature. For those objects, the Newtonian analysis is fine. But light moves at the speed of light, and it therefore sees equal amounts of space and time curvature; so its trajectory bends twice as much as the Newtonian theory predicts.

5.4.5 Numbers!

At the surface of the Earth, the dimensionless gravitational strength is

$$\frac{Gm}{rc^2} \sim \frac{6.7\cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 6.0\cdot 10^{24}\,\mathrm{kg}}{6.4\cdot 10^{6}\,\mathrm{m} \times 3.0\cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0\cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 10^{-9}.$$

This miniscule value is the bending angle (in radians). If physicists want to show that light bends, they had better look beyond the earth! That statement is based on another piece of dimensional analysis and physical reasoning, whose result is quoted without proof: A telescope with mirror of diameter d can resolve angles roughly as small as λ/d, where λ is the wavelength of light; this result is based on the same physics as the diffraction pattern on a CDROM (Section 1.1). One way to measure the bending of light is to measure the change in position of the stars. A lens that could resolve an angle of 10^{-9} has a diameter of at least

$$d \sim \lambda/\theta \sim \frac{0.5\cdot 10^{-6}\,\mathrm{m}}{10^{-9}} \sim 500\,\mathrm{m}.$$

Large lenses warp and crack; one of the largest existing lenses has $d \sim 6\,\mathrm{m}$. No practical mirror can have $d \sim 500\,\mathrm{m}$, and there is no chance of detecting a deflection angle of 10^{-9}.

Physicists therefore searched for another source of light bending. In the solar system, the largest mass is the sun. At the surface of the sun, the field strength is

$$\frac{Gm}{rc^2} \sim \frac{6.7\cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 2.0\cdot 10^{30}\,\mathrm{kg}}{7.0\cdot 10^{8}\,\mathrm{m} \times 3.0\cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0\cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 2.1\cdot 10^{-6} \approx 0.4''.$$

This angle, though small, is possible to detect: The required lens diameter is roughly

$$d \sim \lambda/\theta \sim \frac{0.5\cdot 10^{-6}\,\mathrm{m}}{2.1\cdot 10^{-6}} \sim 20\,\mathrm{cm}.$$

The eclipse expedition of 1919, led by Arthur Eddington of Cambridge University, tried to measure exactly this effect.

For many years – between 1909 and 1916 – Einstein believed that a correct theory of gravity would predict the Newtonian value, which turns out to be 0.87 arcseconds for light that grazes the surface of the sun. The German mathematician Soldner derived the same result in 1803. Fortunately for Einstein's reputation, the eclipse expeditions that went to test his (and Soldner's) prediction got rained or clouded out. By the time an expedition got lucky with the weather (Eddington's in 1919), Einstein had invented a new theory of gravity – general relativity – and it predicted a deflection of 1.75 arcseconds.

The goal of Eddington's expedition was to decide between the Newtonian and general relativity values. The measurements are difficult, and the results were not accurate enough to decide which theory was right. But 1919 was the first year after the World War in which Germany and Britain had fought each other almost to oblivion. A theory invented by a German, confirmed by an Englishman (from Newton's university, no less) – such a picture was comforting after the trauma of war. The world press and scientific community saw what they wanted to: Einstein vindicated!

A proper confirmation of Einstein's prediction came only with the advent of radio astronomy, in whichsmall deflections could be measured accurately. Here is then a puzzle: If the accuracy (resolving power) of a telescope is λ/d, where λ is the wavelength and d is the telescope's diamater, how could radio telescopes be more accurate than optical ones, since radio waves have a much longer wavelength than light?

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of **independent** dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions.}$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{\text{drag}} \sim \rho v^2 A, \tag{5.1}$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

▶ *On what quantities does the drag depend, and what are their dimensions?*

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	$\mathrm{L}^2\mathrm{T}^{-1}$

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

▶ *What dimensionless groups can be constructed for the drag problem?*

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

▶ *Which dimensionless group belongs on the left side?*

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

$$\frac{F}{\rho v^2 r^2} = f\left(\frac{rv}{\nu}\right). \tag{5.3}$$

The physics of the (steady-state) drag force on the cone is all contained in the dimensionless function f. However, dimensional analysis cannot tell us anything about that function. To make progress requires incorporating new knowledge. It can come from an experiment such as dropping the small and large cones, from wind-tunnel tests at various speeds, and even from physical reasoning.

Before reexamining the results of the cone experiment in dimensionless form, let's name the two dimensionless groups. The first one, $F/\rho v^2 r^2$, is traditionally written in a slightly different form:

$$\frac{F}{\frac{1}{2}\rho v^2 A}, \tag{5.4}$$

where A is the cross-sectional area of the cone. The $1/2$ is an arbitrary choice, but it is the usual choice: It is convenient and is reminiscent of the $1/2$ in the kinetic energy formula $mv^2/2$. Written in that way, the first dimensionless group is called the drag coefficient and is abbreviated c_d. The second group, rv/ν, is called the Reynolds number. It is traditionally written in terms of the diameter rather than the radius:

$$\frac{vL}{\nu}, \tag{5.5}$$

where L is the diameter of the object.

The conclusion of the dimensional analysis is then

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{5.6}$$

Now let's see how the cone experiment fits into this dimensionless framework. The experimental data was that the small and large cones fell at the same speed – roughly $1\,\mathrm{m\,s^{-1}}$. The conclusion is that the drag force is proportional to the cross-sectional area A. Because the drag coefficient is proportional to F/A, which is the same for the small and large cones, the small and large cones have the same drag coefficient.

Their Reynolds numbers, however, are not the same. For the small cone, the diameter is $2\,\mathrm{in} \times 0.75$ (why?), which is roughly $4\,\mathrm{cm}$. The Reynolds number is

$$\mathrm{Re} \sim \frac{1\,\mathrm{m\,s^{-1}} \times 0.04\,\mathrm{m}}{1.5 \cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}}, \tag{5.7}$$

where $1\,\mathrm{m\,s^{-1}}$ is the fall speed and $1.5{\cdot}10^{-5}\,\mathrm{m^2\,s^{-1}}$ is the kinematic viscosity of air. Numerically, $\mathrm{Re_{small}} \sim 2000$. For the large cone, the fall speed and viscosity are the same as for the small cone, but the diameter is twice as large, so $\mathrm{Re_{large}} \sim 4000$. The result of the cone experiment is, in dimensionless form, that the drag coefficient is independent of Reynolds number – at least, for Reynolds numbers between 2000 and 4000.

This conclusion is valid for diverse shapes. The most extensive data on drag coefficient versus Reynolds number is for a sphere. That data is plotted logarithmically below (from *Fluid-dynamic Drag: Practical Information on Aerodynamic Drag and Hydrodynamic Resistance* by Sighard F. Hoerner):

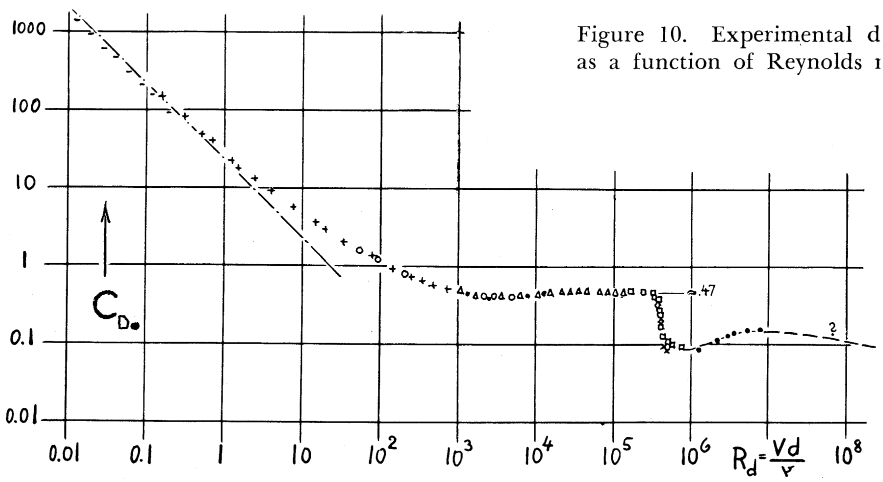

Just like the cones, the sphere's drag coefficient is almost constant in the Reynolds number range 2000 to 4000. This full graph has interesting features. First, toward the low-Reynolds-number end, the drag coefficient increases. Second, for high Reynolds numbers, the drag coefficient stays roughly constant until $\mathrm{Re} \sim 10^6$, where it rapidly drops by almost a factor of 5. The behavior at low Reynolds number will be explained in the chapter on easy (extreme) cases (Chapter 6). The drop in the drag coefficient, which relates to why golf balls have dimples, will be explained in the chapter on lumping (Chapter 7).

Problem 5.4 Only two groups
Show that F, v, r, ρ, and ν produce only two independent dimensionless groups.

Problem 5.5 Counting dimensionless groups
How many independent dimensionless groups are there in the following sets of variables:

a. size of hydrogen including relativistic effects:

$$e^2/4\pi\epsilon_0,\ \hbar,\ c,\ a_0 \text{ (Bohr radius)},\ m_e \text{ (electron mass)}.$$

b. period of a spring–mass system in a gravitational field:

$$T \text{ (period)},\ k \text{ (spring constant)},\ m,\ x_0 \text{ (amplitude)},\ g.$$

c. speed at which a free-falling object hits the ground:

$$v,\ g,\ h \text{ (initial drop height)}.$$

d. [tricky!] weight W of an object:

$$W,\ g,\ m.$$

Problem 5.6 Integrals by dimensions
You can use dimensions to do integrals. As an example, try this integral:

$$I(\beta) = \int_{-\infty}^{\infty} e^{-\beta x^2}\, dx.$$

Which choice has correct dimensions: (a.) $\sqrt{\pi}\beta^{-1}$ (b.) $\sqrt{\pi}\beta^{-1/2}$ (c.) $\sqrt{\pi}\beta^{1/2}$ (d.) $\sqrt{\pi}\beta^{1}$

Hints:

1. The dimensions of dx are the same as the dimensions of x.
2. Pick interesting dimensions for x, such as length. (If x is dimensionless then you cannot use dimensional analysis on the integral.)

Problem 5.7 How to avoid remembering lots of constants
Many atomic problems, such as the size or binding energy of hydrogen, end up in expressions with $\hbar$, the electron mass m_e, and $e^2/4\pi\epsilon_0$, which is a nicer way to express the squared electron charge. You can avoid having to remember those constants if instead you remember these values instead:

$$\hbar c \approx 200\,\text{eV nm} = 2000\,\text{eV Å}$$
$$m_e c^2 \sim 0.5 \cdot 10^6\,\text{eV}$$
$$\frac{e^2/4\pi\epsilon_0}{\hbar c} \equiv \alpha \approx \frac{1}{137} \quad \text{(fine-structure constant)}.$$

Use those values to evaluate the Bohr radius in angstroms (1 Å = 0.1 nm):

$$a_0 = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

As an example calculation using the $\hbar c$ value, here is the energy of a photon:

$$E = hf = 2\pi\hbar f = 2\pi\hbar\frac{c}{\lambda},$$

where f is its frequency and λ is its wavelength. For green light, $\lambda \sim 600\,\text{nm}$, so

$$E \sim \frac{\overbrace{6}^{2\pi} \times \overbrace{200\,\text{eV nm}}^{\hbar c}}{\underbrace{600\,\text{nm}}_{\lambda}} \sim 2\,\text{eV}.$$

Problem 5.8 Heavy nuclei

In lecture we analyzed hydrogen, which is one electron bound to one proton. In this problem you study the innermost electron in an atom such as uranium that has many protons, and analyze one physical consequence of its binding energy.

So, imagine a nucleus with Z protons around which orbits one electron. Let $E(Z)$ be the binding energy (the hydrogen energy is the case $Z = 1$).

a. Show that the ratio $E(Z)/E(1)$ is Z^2.

b. In lecture, we derived that $E(1)$ is the kinetic energy of an electron moving with speed αc where α is the fine-structure constant (roughly 10^{-2}). How fast does the innermost electron move around a heavy nucleus with charge Z?

c. When that speed is comparable to the speed of light, the electron has a kinetic energy comparable to its (relativistic) rest energy. One consequence of such a high kinetic energy is that the electron has enough kinetic energy to produce a positron (an anti-electron) out of nowhere ('pair creation'). That positron leaves the nucleus, turning a proton into a neutron as it exits. So the atomic number Z drops by one: The nucleus is unstable! Relativity sets an upper limit for Z.

   Estimate that maximum Z and compare it with the Z for the heaviest stable nucleus (uranium).

Problem 5.9 Power radiated by an accelerating charge

Electromagnetism, where the usual derivations are so cumbersome, is an excellent area to apply dimensional analysis. In this problem you work out the power radiated by an accelerating charge, which is how radio stations work.

So, consider a particle with charge q, with position x, velocity v, and acceleration a. What variables are relevant to the radiated power P? The position cannot matter because it depends on the origin of the coordinate system, whereas the power radiated cannot depend on the origin. The velocity cannot matter because of relativity: You can transform to a reference frame where $v = 0$, but that change will not affect the radiation (otherwise you could distinguish a moving frame from a non-moving frame, in violation of the principle of relativity). So the acceleration a is all that's left to determine the radiated power. [This line of argument is slightly dodgy, but it works for low speeds.]

a. Using P, $q^2/4\pi\epsilon_0$, and a, how many dimensionless quantities can you form?

b. Fix the problem in the previous part by adding one quantity to the list of variables, and give a physical reason for including the quantity.

c. With the new list, use dimensionless groups to find the power radiated by an accelerating point charge. In case you are curious, the exact result contains a dimensionless factor of 2/3; dimensional analysis triumphs again!

Problem 5.10 Yield from an atomic bomb

Geoffrey Taylor, a famous Cambridge fluid mechanic, annoyed the US government by doing the following analysis. The question he answered: 'What was the yield (in kilotons of TNT) of the first atomic blast (in the New Mexico desert in 1945)?' Declassified pictures, which even had a scale bar, gave the following data on the radius of the explosion at various times:

t (ms)	R (m)
3.26	59.0
4.61	67.3
15.0	106.5
62.0	185.0

a. Use dimensional analysis to work out the relation between radius R, time t, blast energy E, and air density ρ.

b. Use the data in the table to estimate the blast energy E (in Joules).

c. Convert that energy to kilotons of TNT. One gram of TNT releases 1 kcal or roughly 4 kJ.

The actual value was 20 kilotons, a classified number when Taylor published his result ['The Formation of a Blast Wave by a Very Intense Explosion. II. The Atomic Explosion of 1945.', *Proceedings of the Royal Society of London. Series A, Mathematical and Physical* **201**(1065): 175–186 (22 March 1950)]

Problem 5.11 Atomic blast: A physical interpretation

Use energy densities and sound speeds to make a rough physical explanation of the result in the 'yield from an atomic bomb' problem.

Problem 5.12 Rolling down the plane

Four objects, made of identical steel, roll down an inclined plane without slipping. The objects are:

1. a large spherical shell,
2. a large disc,
3. a small solid sphere,
4. a small ring.

The large objects have three times the radius of the small objects. Rank the objects by their acceleration (highest acceleration first).

Check your results with exact calculation or with a home experiment.

Problem 5.13 Blackbody radiation

A hot object – a so-called blackbody – radiates energy, and the flux F depends on the temperature T. In this problem you derive the connection using dimensional analysis. The goal is to find F as a function of T. But you need more quantities.

a. What are the dimensions of flux?

b. What two constants of nature should be included because blackbody radiation depends on the quantum theory of radiation?

c. What constant of nature should be included because you are dealing with temperature?

d. After doing the preceding parts, you have five variables. Explain why these five variables produce one dimensionless group, and use that fact to deduce the relation between flux and temperature.

e. Look up the Stefan–Boltzmann law and compare your result to it.

Part 3

Lossy compression

6
Easy cases

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

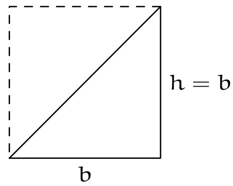

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what b and h are*, so $A = bh/2$.

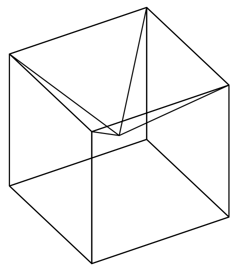

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

6.2 Atwood machine

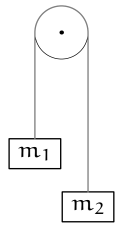

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

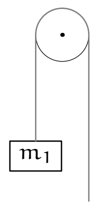

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

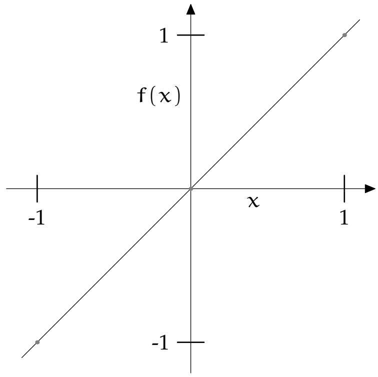

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

Look how simple the result is when derived in a symmetric, dimensionless form using easy cases.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{6.1}$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 31].)

In other words, f is constant! The consequence is

$$F_{\mathrm{d}} \sim \rho v^2 A, \tag{6.2}$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_{\mathrm{d}} \equiv \frac{F_{\mathrm{d}}}{\rho v^2 A} \tag{6.3}$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_{d} for various shapes (at high Reynolds number).

Object	c_{d}
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [21], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re $\ll$ 1):

$$c_{\mathrm{d}} = f(\mathrm{Re}) \qquad (\text{for } \mathrm{Re} \ll 1). \tag{6.4}$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \tag{6.5}$$

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_{\mathrm{d}} \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_{\mathrm{d}}}{\rho_{\mathrm{fl}} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_{d} proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $\mathrm{Re} \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_{\mathrm{d}} \sim \rho_{\mathrm{fl}} r^2 v^2 \frac{\nu}{vr} = \rho_{\mathrm{fl}} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_{\mathrm{d}} = 6\pi \rho_{\mathrm{fl}} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

6.3.3 Conductivity of seawater

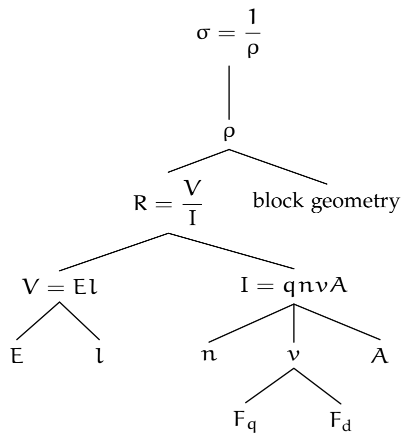

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_{d}. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

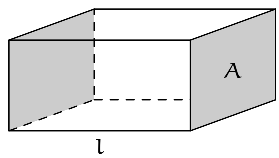

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_{\mathrm{d}} = 6\pi\rho_{\mathrm{fl}}\nu v r. \qquad (6.6)$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will contain n.) To find σ climb the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2nAV}{6\pi\eta lr}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta lr}{q^2nA}.$$

The voltage V has vanished, which is encouraging: In most circuits, the conductivity (and resistance) is independent of voltage. Use the resistance to find the resistivity:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2n}.$$

The expressions simplify as we climb the tree. For example, the geometric parameters l and A have vanished. Their disappearance is encouraging – the purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Now use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier. Thus,

$$\sigma = \frac{1}{\rho} \sim \frac{e^2n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. To do so, try a kitchen-sink experiment. Add table salt to a glass of water until it tastes as salty as seawater. I just tried it. In a glass of water, I found that one teaspoon of salt tastes like drinking seawater. A glass of water may have a volume of $0.3\,\ell$ or a mass of 300 g. A flat teaspoon of salt has a volume of about $5\,\mathrm{m}\ell$. Why $5\,\mathrm{m}\ell$? A teaspoon is about 4 cm long by 2 cm wide by 1 cm thick at its deepest point; let's assume 0.5 cm on average. Its volume is therefore

$$\text{teaspoon} \sim 4\,\mathrm{cm} \times 2\,\mathrm{cm} \times 0.5\,\mathrm{cm} \sim 4\,\mathrm{cm}^3.$$

The density of salt is maybe twice the density of water, so a flat teaspoon has a mass of roughly 10 g. The mass fraction of salt in seawater is, in this experiment, roughly 1/30. The true value is remarkably close: 0.035. A mole of salt, which provides two charges per NaCl 'molecule', has a mass of 60 g, so

$$n \sim \frac{1}{30} \times \underbrace{1\,\mathrm{g\,cm^{-3}}}_{\rho_{\mathrm{water}}} \times \frac{2\,\text{charges}}{\text{molecule}} \times \frac{6\cdot 10^{23}\,\text{molecules mole}^{-1}}{60\,\mathrm{g\,mole^{-1}}}$$
$$\sim 7\cdot 10^{20}\,\text{charges cm}^{-3}.$$

With n evaluated, the only remaining mysteries in the conductivity

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}$$

are the ion radius r and the dynamic viscosity η.

Do the easy part first. The dynamic viscosity is

$$\eta = \rho_{\mathrm{water}}\nu \sim 10^3\,\mathrm{kg\,m^{-3}} \times 10^{-6}\,\mathrm{m^2\,s^{-1}} = 10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units than in SI units, the one exception is electromagnetism, which is represented by the e^2 in the conductivity. Electromagnetism is conceptually easier in cgs units – which needs no factors of μ_0 or $4\pi\epsilon_0$, for example – than in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations in electromagnetism, use SI units.

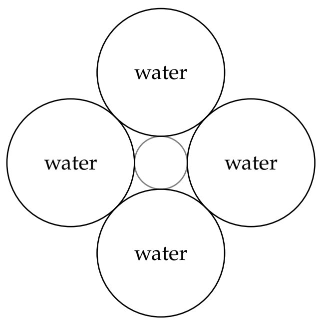

The final quantity required is the ion radius. A positive ion (sodium) attracts an oxygen end of a water molecule; a negative ion (chloride) attracts the hydrogen end of a water molecule. Either way, the ion, being charged, is surrounded by one or maybe more layers of water molecules. As it moves, it drags some of this baggage with it. So rather than use the bare ion radius you should use a larger radius to include this shell. But how thick is the shell? As a guess, assume that the shell includes one layer of water molecules, each with a radius of 1.5 Å. So for the ion plus shell, $r \sim 2$ Å.

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6\cdot 10^{-19}\,\mathrm{C})^2}^{e^2}\times\overbrace{7\cdot 10^{26}\,\mathrm{m}^{-3}}^{n}}{\underbrace{6\times 3}_{6\pi}\times\underbrace{10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}}_{\eta}\times\underbrace{2\cdot 10^{-10}\,\mathrm{m}}_{r}}.$$

You can do the computation mentally. First count the powers of ten and then worrying about the small factors. Then count the top and bottom contributions separately. The top contributes −12 powers of 10: −38 from e^2 and +26 from n. The bottom contributes −13 powers of 10: −3 from η and −10 from r. The division produces one power of 10.

Now account for the remaining small factors:

$$\frac{1.6^2\times 7}{6\times 3\times 2}.$$

Slightly overestimate the answer by pretending that the 1.6^2 on top cancels the 3 on the bottom. Slightly underestimate the answer – and maybe compensate for the overestimate – by pretending that the 7 on top cancels the 6 on the bottom. After these lies, only 1/2 remains. Multiplying it by the sole power of ten gives

$$\sigma \sim 5\,\Omega^{-1}\,\mathrm{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\mathrm{m}^{-1}$, which is extremely close to the result from mental calculation.

The estimated resistivity is

$$\rho \sim \sigma^{-1} \sim 0.2\,\Omega\,\mathrm{m} = 20\,\Omega\,\mathrm{cm},$$

where we converted to the conventional although not fully SI units of $\Omega\,\mathrm{cm}$. A typical experimental value for seawater at $T = 15\,^\circ\mathrm{C}$ is $23.3\,\Omega\,\mathrm{cm}$ (from [15, p. **14-**15]), absurdly close to the estimate!

Probably the most significant error is the radius of the ion-plus-water combination that is doing the charge transport. Perhaps r should be greater than 2 Å, especially for a sodium ion, which is smaller than chloride; it therefore has a higher electric field at its surface and grabs water molecules more strongly than chloride does. In spite of such uncertainties, the continuum approximation produced accurate results.

That accuracy is puzzling. At the length scale of a sodium ion, water looks like a collection of spongy boulders more than it looks like a continuum. Yet Stokes drag worked. It works because the important length scale is not the size of water molecules, but rather their mean free path between collisions. Molecules in a liquid are packed to the point of contact, so the mean free path is much shorter than a molecular (or even ionic) radius, especially compared to an ion with its shell of water.

The moral of this example, besides the application of Stokes drag, is to have courage: Approximate first and ask questions later. The approximations might be accurate for reasons that you do not suspect when you start solving a problem. If you agonize over each approximation, you will never start a calculation, and then you will not find out that many approximations would have been fine – if only you had had the courage to make them.

6.3.4 Combining solutions from the two limits

You know know the drag force in two extreme cases, viscous and turbulent drag. The results are repeated here:

$$F_{\mathrm{d}} = \begin{cases} 6\pi\rho_{\mathrm{fl}}\nu v r & \text{(viscous)}, \\ \frac{1}{2}c_{\mathrm{d}}\rho_{\mathrm{fl}}Av^2 & \text{(turbulent)}. \end{cases}$$

Let's compare and combine them by making the viscous form look like the turbulent form: Multiply by the Reynolds number rv/ν (basing the Reynolds number on radius rather than diameter). Then

$$F_{\mathrm{d}} = \underbrace{\left(\frac{rv/\nu}{\mathrm{Re}}\right)}_{1} \times \underbrace{6\pi\rho_{\mathrm{fl}}v\nu r}_{F_{\mathrm{d}}} = \frac{1}{\mathrm{Re}}\,6\pi\rho_{\mathrm{fl}}v^2r^2 \qquad \text{(viscous)}.$$

With $A = \pi r^2$,

$$F_{\mathrm{d}} = \frac{6}{\mathrm{Re}}\,\rho_{\mathrm{fl}}v^2A \qquad \text{(viscous)},$$

Since $c_{\mathrm{d}} \equiv F_{\mathrm{d}}/(\frac{1}{2}\rho_{\mathrm{fl}}v^2A)$,

$$c_{\mathrm{d}} = \frac{12}{\mathrm{Re}} \qquad \text{(viscous)}.$$

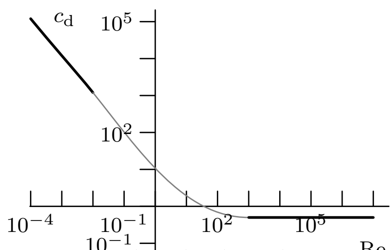

That limit and the the high-speed limit $c_{\mathrm{d}} \sim 0.5$ are sketched in the graph (with a gray interpolation between the limits). Almost all of the experimental data is explained by this graph, except for the drop in c_{d} near the Re $\sim 10^6$.

More on special cases

Summary:

Further reading:

Problem 6.1 Integrals
Use special cases of a to choose the correct formula for each integral.

a. $\int_{-\infty}^{\infty} e^{-ax^2}\,dx$

(1.) $\sqrt{\pi a}$ (2.) $\sqrt{\pi/a}$

b. $\int_{-\infty}^{\infty} \frac{1}{a^2+x^2}\,dx$

(1.) πa (2.) π/a (3.) $\sqrt{\pi}a$ (4.) $\sqrt{\pi}/a$

Problem 6.2 Debugging
Use special (i.e. easy) cases of n to decide which of these two C functions correctly computes the sum of the first n odd numbers:

```
int sum_of_odds (int n) {
  int i, total = 0;
  for (i=1; i<=2*n+1; i+=2)
    total += i;
  return total;
}
```

or

```
int sum_of_odds (int n) {
  int i, total = 0;
  for (i=1; i<=2*n-1; i+=2)
    total += i;
  return total;
}
```

Problem 6.3 Reynolds numbers
Estimate the Reynolds number for:

a. a falling raindrop;

b. a flying mosquito;

Problem 6.4 Drag at low Reynolds number
At low Reynolds number, the drag on a sphere is

$$F = 6\pi\rho\nu v r.$$

What is the drag coefficient c_d as a function of Reynolds number Re?

Problem 6.5 Truncated pyramid

In this problem you use special cases to find the volume of a truncated pyramid. It has a square base with side b, a square top with side a, and height h. So, use special cases of a and b to evaluate these candidates for the volume:

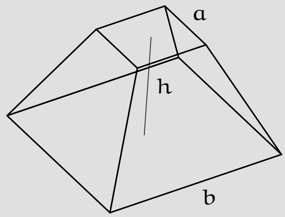

a. $\frac{1}{3}hb^2$

b. $\frac{1}{3}ha^2$

c. $\frac{1}{3}h(a^2 + b^2)$

d. $\frac{1}{2}h(a^2 + b^2)$

Which if any of these formulas pass all your special-cases tests? If no formula passes all tests, invent a formula that does. If you are stuck, find the volume by integration!

Problem 6.6 Fog

a. Estimate the terminal speed of fog droplets (radius $\sim 10\,\mu$m). Use either the low- or high-Reynolds-number limit for the drag force, whichever you guess is the more likely to be valid.

b. Use the speed to estimate the Reynolds number and check that you used the correct limit for the drag force. If not, try the other limit!

   It is much less than 1, so the original assumption of low-Reynolds-number flow is okay.

c. Fog is a low-lying cloud. How long would fog droplets take to fall 1 km (the height of a typical cloud)? What is the everyday effect of this settling time?

Problem 6.7 Tube flow

In this problem you study fluid flow through a narrow tube. The quantity to predict is Q, the volume flow rate (volume per time). This rate depends on:

l	the length of the tube
Δp	the pressure difference between the tube ends
r	the radius of the tube
ρ	the density of the fluid
ν	the kinematic viscosity of the fluid

a. Find three independent dimensionless groups G_1, G_2, and G_3 from these six variables. *Hint 1:* One physically reasonable group is $G_2 = r/l$. *Hint 2:* Put Q in G_1 only! Then write the general form

$$G_1 = f(G_2, G_3).$$

[There are lots of choices for G_1 and G_3.]

b. Now imagine that the tube is very long and thin ($l \gg r$) and that the radius or flow speed are small enough to make the Reynolds number low. Then you can deduce the form of f using proportional reasoning.

You might think about these proportionalities:

1. How should Q depend on Δp? For example, if you double the pressure difference, what should happen to the flow rate?
2. How should Q depend on l? For example, if you keep the pressure difference the same but double the tube length, what happens to Q? Or if you double Δp and double l, what happens to Q?

Figure out the form of f to satisfy all your proportionality requirements.

If you get stuck going forward, instead work backward from the correct result. Look up Poiseuille flow, and use this result to deduce the preceding proportionalities; and then give reasons for why they are that way.

c. *[optional]*
The dimensional analysis in the preceding parts does not tell you the dimensionless constant. Use a syringe and needle to estimate the constant. Compare your constant with the value of $\pi/8$ that comes from solving the equations of fluid mechanics honestly.

Problem 6.8 Atwood machine: Tension in the string

Here is the Atwood machine from lecture. The string and pulley are massless and frictionless. We used dimensional analysis and special cases to guess the acceleration of either mass. With the right choice of sign,

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

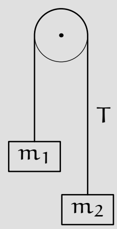

In this problem you guess the tension in the string.

a. The tension T, like the acceleration, depends on m_1, m_2, and g. Explain why these four variables result in two independent dimensionless groups.

b. Choose two suitable independent dimensionless groups so that you can write an equation for the tension in this form:

dimensionless group containing $T = f$(dimensionless group without T).

The next part will be easier if you use a lot of symmetry in choosing the groups.

c. Use special cases to guess f, and sketch f.

d. Solve for T using the usual methods from introductory physics (8.01); then compare that answer with your answer from the preceding part.

Problem 6.9 Plant-watering system

The semester is over and you are going on holiday for a few weeks. But how will you water the house plants?! Design an unpowered slow-flow system to keep your plants happy.

Problem 6.10 Dimensional analysis for circuits

a. Using Q as the dimension of charge, what are the dimensions of inductance L, capacitance C, and resistance R?

b. Show that the dimensions of L, C, and R contain two independent dimensions.

c. In a circuit with one inductor, one capacitor, and one resistor, one dimensionless group should result from the three component values L, R, and C. What is physical interpretation of this group?

Problem 6.11 Tipsy host for the three-doors problem

7 Lumping

The symmetry chapter (Section 3.1) introduced the principle of invariance: ‘When there is change, look for what does not change.’ However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, ‘Make it so.’

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

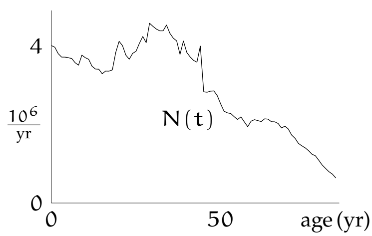

$$N_{\text{babies}} = \int_0^{2\,\text{yr}} N(t)\,dt. \tag{7.1}$$

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

▶ *Why did the life expectancy drop from* 80 *to* 75 *years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

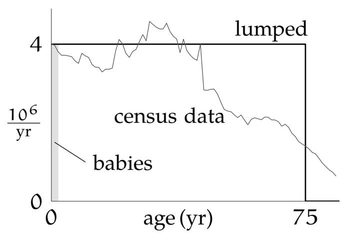

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

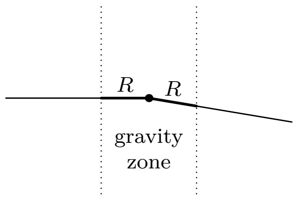

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

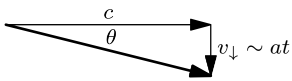

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

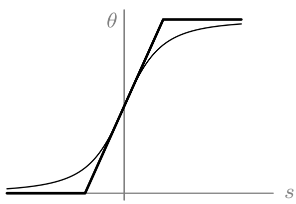

...the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

Problem 7.4 Higher values of GM/Rc^2
When GM/Rc^2 is no longer small, strange things happen. Use lumping to predict what happens to light when $GM/Rc^2 \sim 1$.

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

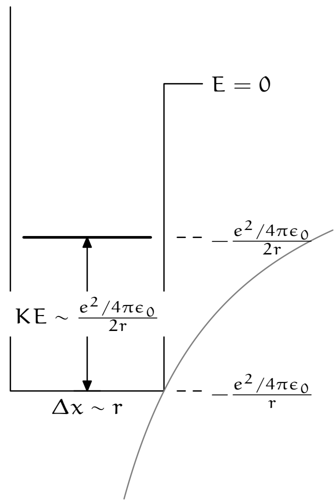

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{7.5}$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

There are two approximate methods to determine the minimum r. The first method is familiar from the analysis of lift in Section 3.5.2: When two terms compete, the minimum occurs when the terms are roughly equal. In other words, at the minimum energy, the potential energy and kinetic energy (the two competing terms) are roughly equal in magnitude. Using the Bohr radius a_0 as the corresponding separation, this criterion says

$$\underbrace{\frac{e^2}{4\pi\epsilon_0 a_0}}_{\text{PE}} \sim \underbrace{\frac{\hbar^2}{m_e a_0^2}}_{\text{KE}}. \tag{7.6}$$

The result for a_0 is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4/4\pi\epsilon_0)}. \tag{7.7}$$

The second method of estimating the minimum-energy separation is to use dimensional analysis, by writing the energy and radius in dimensionless forms. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is therefore worth learning via this example. To make r dimensionless, cook up another length l and then define $\bar{r} \equiv r/l$. The only other length that is based on the parameters of hydrogen (and the relevant constants of nature such as ϵ_0) is

$$l \equiv \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

So, let's defined the scaled (dimensionless) radius as

$$\bar{r} \equiv \frac{r}{l}.$$

To make the energy dimensionless, cook up another energy based on the parameters of the problem. A reasonable candidate for this energy scale is $e^2/4\pi\epsilon_0 l$. That choice defines the scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the scaled length and energy, the total energy

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

simplifies greatly:

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants all live in the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

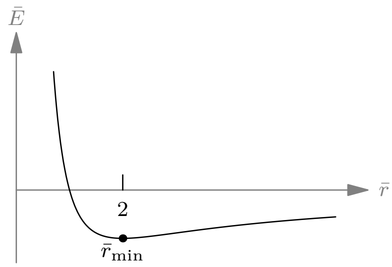

Calculus locates this minimum-energy $\bar{r}$ at $\bar{r}_{\min} = 2$. Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{\min} \sim 1$. In normal, unscaled terms, it is

$$r_{\min} = l\bar{r}_{\min} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius as computed using dimensional analysis (Section 5.3.1) and is also the exact Bohr radius computed properly using quantum mechanics. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization! Even if the method were not so charmed, there is no point in doing a proper calculus minimization. Given the inaccuracies in the rest of the derivation, The calculus method is too accurate

Engineers understand this idea of not over-engineering a system. If a bicycle most often breaks at welds in the frame, there is little point replacing the metal between the welds with expensive, high-strength aerospace materials. The new materials might last 100 years, but such a replacement would be overengineering because something else will break before 100 years are done.

In estimating the Bohr radius, the kinetic-energy estimate uses a crude form of the uncertainty principle, $\Delta p \Delta x \sim \hbar$, whereas the true statement is that $\Delta p \Delta x \geq \hbar/2$. The estimate also uses the approximation $E_{\text{kinetic}} \sim (\Delta p)^2/m$. This approximation contains m instead of $2m$ in the denominator. It also assumes that Δp can be converted into an energy as though it were a true momentum rather than merely a crude estimate for the root-mean-square momentum. The potential- and kinetic-energy estimates use a crude definition of position uncertainty Δx: that $\Delta x \sim r$. After making so many approximations, it is pointless to minimize the result using the elephant gun of differential calculus. The approximate method is as accurate as, or perhaps more accurate than the approximations in the energy.

The method of equating competing terms is called balancing. We balanced the kinetic energy against the potential energy by assuming that they are roughly the same size. Nature could have been unkind: The potential and kinetic energies could have differed by a factor of 10 or 100. But Nature is kind: The two energies are roughly equal, except for a constant that is nearly 1 (of order unity). This rough equality occurs in many examples: You often get a reasonable answer simply by pretending that two energies (or two quantities with the same units) are equal. [When the quantities are potential and kinetic energy, as they often are, you get extra safety: The so-called virial theorem protects you against large errors (for more on the virial theorem, see any intermediate textbook on classical dynamics).]

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_{d} is a function only of the Reynolds number:

$$\frac{F_{\mathrm{d}}}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_{\mathrm{d}}/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

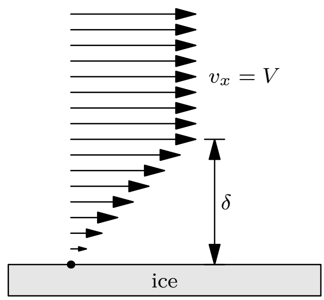

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are $\mathrm{L}^2\mathrm{T}^{-1}$. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \tag{7.10}$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \tag{7.11}$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \tag{7.12}$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

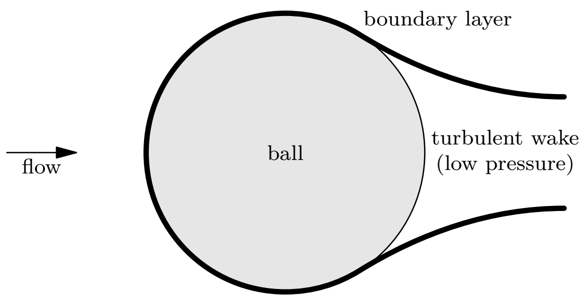

That interpretation of the drag coefficient leads to an explanation of why the drag coefficient remains roughly constant as the Reynolds number goes to infinity – in other words, as the flow speed increases or the viscosity decreases. In that limit, the boundary-layer detachment point shifts as far forward as it goes, namely to the widest portion of the object. Then the drag coefficient is roughly 1.

This explanation mostly accounts for the high-Re data on drag coefficient versus Reynolds number. Here is the log–log plot from Section 5.6:

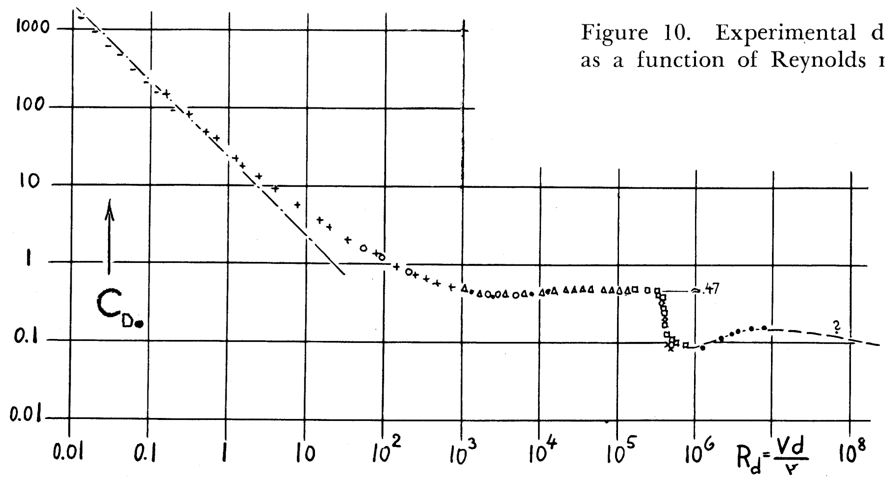

Figure 10. Experimental d
as a function of Reynolds ı

The drag coefficient is roughly constant in the Reynolds-number range 10^3 to (almost) 10^6. But why does the drag coefficient drop sharply around $\mathrm{Re} \sim 10^6$? The boundary-layer picture can also help us understand this behavior. To do so, first compute Re_δ, the Reynolds number of the flow in the boundary layer. The Reynolds number is defined as

$$\mathrm{Re} = \frac{\text{typical flow speed} \times \text{distance over which the flow speed varies}}{\text{kinematic viscosity}}. \tag{7.13}$$

In the boundary layer, the flow speed varies from 0 to v, so it is comparable to v. The speed changes over the boundary-layer thickness δ. So

$$\mathrm{Re}_\delta \sim \frac{v\delta}{\nu}. \tag{7.14}$$

Because $\delta \sim r \times \mathrm{Re}^{-1/2}$,

$$\mathrm{Re}_\delta \sim \frac{vr}{\nu} \times \mathrm{Re}^{-1/2}. \tag{7.15}$$

The first fraction is just the regular Reynolds number Re, so

$$\mathrm{Re}_\delta \sim \mathrm{Re} \times \mathrm{Re}^{-1/2} = \mathrm{Re}^{1/2}. \tag{7.16}$$

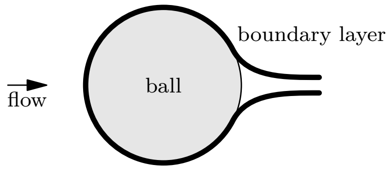

Flows become turbulent when $\mathrm{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\mathrm{Re}_\delta \sim 10^3$ or $\mathrm{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

Problem 7.5 Perfume

If the diffusion constant (in air) for small perfume molecules is $10^{-6}\,\mathrm{m^2\,s^{-1}}$, estimate the time for perfume molecules to diffuse across a room.

Now try the experiment: How long does it take to smell the perfume from across the room? Explain the large discrepancy between the theoretical estimate and the experimental value.

Problem 7.6 Bending of light

In lecture we estimated how much gravity bends light by using dimensional analysis and then guessing the final functional form. In this problem, you analyze the same situation using lumping.

As before, let r be the closest that the light gets to the origin (which is the center of the gravitating mass). Discretize: Pretend that gravity forgets to deflect the photon unless the photon's distance to the origin is comparable to r. Using that idea, estimate the radial (inward) velocity imparted to the photon, and then the bending angle.

Feel free to neglect dimensionless constants like 2 or π, and check your answer against what we derived in lecture.

Problem 7.7 Teacup spindown

You stir your afternoon tea to mix the milk (and sugar if you have a sweet tooth). Once you remove the stirring spoon, the rotation starts to slow. What is the spindown time τ? In other words, how long before the angular velocity of the tea has fallen by a significant fraction?

To estimate τ, consider a physicist's idea of a teacup: a cylinder with height L and diameter L, filled with liquid. Why does the rotation slow? Tea near the edge of the teacup – and near the base, but for simplicity neglect the effect of the base – is slowed by the presence of the edge (the no-slip boundary condition). The edge produces a velocity gradient.

Because of the tea's viscosity, the velocity gradient produces a force on any piece of the edge. This force tries to spin the piece in the direction of the tea's motion. The piece exerts a force on the tea equal in magnitude and opposite in direction. Therefore, the edge slows the rotation. Now you can analyze this model quantitatively.

a. In terms of the total viscous force F and of the initial angular velocity ω, estimate the spindown time. Hint: Consider torque and angular momentum. (Feel free to drop any constants, such as π and 2, by invoking the Estimation Theorem: $1 = 2$.)

b. You can estimate F with the idea that

$$\text{viscous force} \sim \rho\nu \times \text{velocity gradient} \times \text{surface area}.$$

Here $\rho\nu$ is η. The more familiar viscosity is η, known as the dynamic viscosity. The more convenient viscosity is ν, the kinematic viscosity. The velocity gradient is determined by the size of the region in which the the edge has a significant effect on the flow; this region is called the boundary layer. Let δ be its thickness. Estimate the velocity gradient near the edge in terms of δ, and use the equation for viscous force to estimate F.

c. Put your expression for F into your earlier estimate for τ, which should now contain only one quantity that you have not yet estimated (the boundary-layer thickness).

d. You can estimate δ using your knowledge of random walks. The boundary layer is a result of momentum diffusion; just as D is the molecular-diffusion coefficient, ν is the momentum-diffusion coefficient. In a time t, how far can momentum diffuse? This distance is δ. What is a natural estimate for t? (Hint: After rotating 1 radian, the fluid is moving in a significantly different direction than before, so the momentum fluxes no longer add.) Use that time to estimate δ.

e. Now put it all together: What is the characteristic spindown time τ (the time for the rotation to slow down by a significant amount)?

f. Stir some tea to experimentally estimate τ_{exp}. Compare this time with the time predicted by the preceding theory. [In water (and tea is roughly water), $\nu \sim 210^{-6}\,\text{m}^2\,\text{s}^{-1}$.]

Problem 7.8 Stokes' law

You can use ideas from Problem 7.7 to derive Stokes' formula for drag at low speeds (more precisely, at low Reynolds' number). In the text we derived the result from dimensional analysis; here you will develop a physical argument.

Consider a sphere of radius R moving with velocity v. Equivalently, in the reference frame of the sphere, the sphere is fixed and the fluid moves past it with velocity v. Next to the sphere, the fluid is stationary. Over a region of thickness δ (the boundary layer), the fluid velocity rises from zero to the full flow speed v. Assume that $\delta \sim R$ (the most natural assumption) and estimate the viscous drag force. Compare the force with Stokes' formula (remember that $\rho\nu = \eta$).

Problem 7.9 Bouncing ball

You drop a steel ball, say $r \sim 1\,\text{cm}$, from a height of one or two metres. It lands on a steel surface and bounces up to nearly the original drop height. Estimate the contact force at the instant when the ball is stationary on the ground. Give your answer as the dimensionless ratio

$$\frac{\text{contact force}}{\text{weight of the ball}}.$$

Useful data: The elastic modulus of steel is $Y \sim 10^{11}\,\text{N}\,\text{m}^{-2}$.

Problem 7.10 Cone free-fall time and distance

Estimate how long the falling cones of Section 3.3.1 require to reach (a good fraction of) terminal velocity. And estimate how far they fall before reaching (a good fraction of) terminal velocity.

Problem 7.11 Kinetic theory

Molecules in a gas travel in all directions and with a continuous range of speeds. However, many results can be understood with the following lumping approximation: Pretend that all molecules move with the thermal speed and that they move only along a coordinate axis (i.e. in one of six directions).

Using the preceding lumping approximation, compute the following quantities:

Problem 7.12 Electric field of a uniform sheet of charge

Problem 7.13 Electric field inside a uniform shell of charge

8
Probabilistic reasoning

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \qquad (8.2)$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \qquad (8.3)$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \qquad (8.4)$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \qquad (8.5)$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \qquad (8.6)$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19 . . . 21 cm and the length lies in the range 28 . . . 32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of 300 . . . 3000 cm^2 spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300\ldots3000\,\text{cm}^2.$$

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

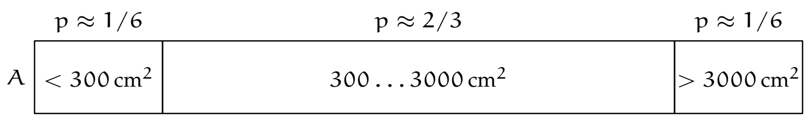

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

28 . . . 32 cm and 19 . . . 21 cm respectively. Here are their probability interpretations:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
h	< 28 cm	28 . . . 32 cm	> 32 cm

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
w	< 19 cm	19 . . . 21 cm	> 21 cm

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is an answer to this question that depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

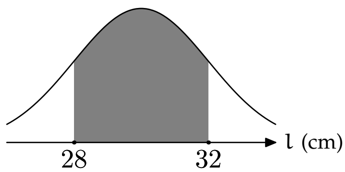

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution, meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu - \sigma$ to $\mu + \sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu - \sigma$ to $\mu + \sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28 \times 32}$ cm or approximately 29.93 cm.

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space.

Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints:

$$\mu = \mu_1 \mu_2. \tag{8.9}$$

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

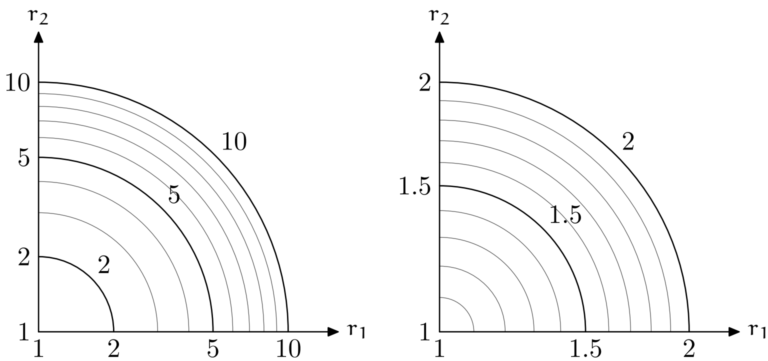

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

Problem 8.1 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example where the width and height ranges are $1 \ldots 2\,\text{m}$. What is the plausible range for the area? The naive approach of simply multiplying endpoints produces a plausible range of $1 \ldots 4\,\mathrm{m}^2$ – a width of a factor of 4. However, this range is too pessimistic and the correct range should be narrower. Using the log-normal distribution, the range spans a factor of

$$\exp\left(\sqrt{2 \times (\ln 2)^2}\right) \approx 2.67.$$

This span and the midpoint determine the range. The area midpoint is the product of the width and height midpoints, each of which is $\sqrt{2}\,\mathrm{m}$. So the midpoint is $2\,\mathrm{m}^2$. The correct endpoints of the plausible range are therefore

$$\frac{2\,\mathrm{m}^2}{\sqrt{2.67}} \ldots 2\,\mathrm{m}^2 \times \sqrt{2.67}$$

or $1.23 \ldots 3.27\,\mathrm{m}^2$. In other words, I assign roughly a 1/6 probability that the area is less than $1.23\,\mathrm{m}^2$ and roughly a 1/6 probability that it is greater than $3.27\,\mathrm{m}^2$. Those conclusions seem reasonable when using such uncertain knowledge of length and width.

Having checked that the method is reasonable, it is time to test it in the original illustrative problem: the plausible area range for an A4 sheet. The naive plausible range was $532 \ldots 672\,\mathrm{cm}^2$, and the correct plausible range will be narrower. Indeed, the log-normal method gives the narrower area range of $550 \ldots 650\,\mathrm{cm}^2$ with a best guess (most likely value) of $598\,\mathrm{cm}^2$. How did we do? The true area is exactly $2^{-4}\,\mathrm{m}^2$ or $625\,\mathrm{cm}^2$ because – I remembered only after doing this calculation! – An paper is constructed to have one-half the area of A$(n-1)$ paper, with A0 paper having an area of $1\,\mathrm{m}^2$. The true area is only 5% larger than the best guess, suggesting that we used accurate information about the length and width; and it falls within the plausible range but not right at the center, suggesting that the method for computing the plausible range is neither too daring nor too conservative.

The analysis of combining ranges illustrates the two crucial points about divide-and-conquer reasoning. First, the main benefit comes from subdividing vague knowledge (such as the area itself) into pieces about which our knowledge is accurate (the length and the width). Second, this benefit swamps the small penalty in accuracy that results from combining many quantities together.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

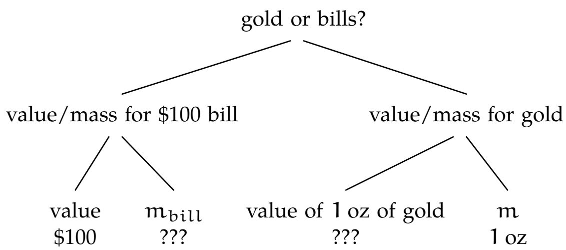

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

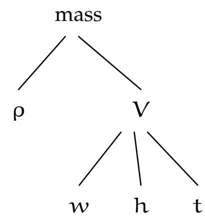

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6\,\text{cm}$ and $w \sim 15\,\text{cm}$. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

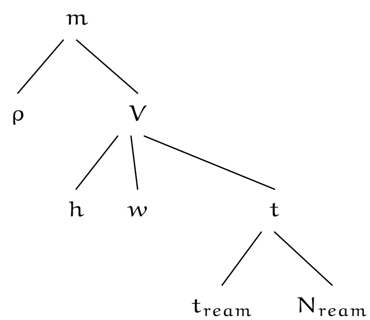

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

▶ *What is the estimate for ρ, the density of a bill?*

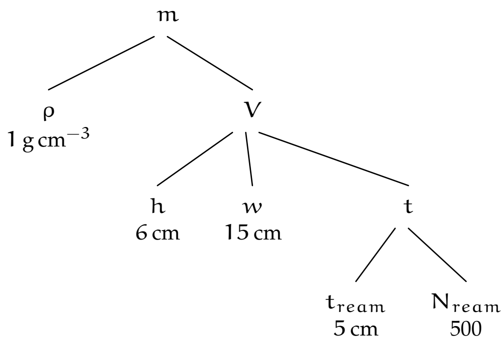

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\text{g}\,\text{cm}^{-3}$. Therefore the density of a bill is roughly $1\,\text{g}\,\text{cm}^{-3}$.

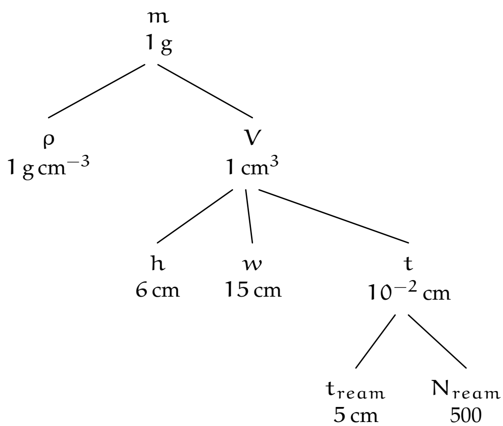

Now propagate the leaf values upward. The thickness of a bill is roughly $10^{-2}\,\text{cm}$, so the volume of a bill is roughly

$$\begin{aligned} V &\sim 6\,\text{cm} \times 15\,\text{cm} \times 10^{-2}\,\text{cm} \\ &\sim 1\,\text{cm}^3. \end{aligned}$$

The mass is therefore

$$m \sim 1\,\text{cm}^3 \times 1\,\text{g}\,\text{cm}^{-3} \sim 1\,\text{g}.$$

and the value per mass of an $\$N$ bill is therefore $\$N/\text{g}$. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of \$800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were \$840/oz, the arithmetic is simple enough to do mentally, and produces \$30/g. An exact division produces the slightly lower figure of \$28/g. The result of this calculation is as follows: In the bank vault, first collect all the \$100 bills that we can carry. If we have spare capacity, collect the \$50 bills, the gold, and only then the \$20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes and propagate them upward – thereby obtaining plausible ranges for the value per mass of bills and gold.

> **Problem 8.2 Your plausible ranges**
> What are your plausible ranges for the five leaf quantities t_{ream}, N_{ream}, w, h, and ρ? Propagate them upward to get plausible ranges for the interior nodes including for the root node m.

Here are my ranges along with a few notes on how I estimated a few of them:

1. thickness of a ream, t_{ream}: $4\ldots6\,\text{cm}$.
2. number of sheets in a ream, N_{ream}: 500. I'm almost certain that I remember this value correctly, but to be certain I confirmed it by looking at a label on a fresh ream.
3. width of a bill, w: $10\ldots20\,\text{cm}$. A reasonable length estimate seemed to be 6 in but I could give or take a couple inches. In metric units, $4\ldots8$ in becomes (roughly) $10\ldots20\,\text{cm}$.
4. height of a bill, h: $5\ldots7\,\text{cm}$.
5. density of a bill, ρ: $0.8\ldots1.2\,\text{g}\,\text{cm}^{-3}$. The argument for $\rho = 1\,\text{g}\,\text{cm}^{-3}$ – that a bill is made from paper and paper is made from wood – seems reasonable. However, the many steps required to process wood into paper may reduce or increase the density slightly.

Now propagate these ranges upward. The plausible range for the thickness t becomes $0.8\ldots1.2\cdot10^{-2}\,\text{cm}$. The plausible range for the volume V becomes $0.53\ldots1.27\,\text{cm}^3$. The plausible range for the mass m becomes $0.50\ldots1.30\,\text{g}$. The plausible range for the value per mass is \$79...189/g (with a midpoint of \$122/g).

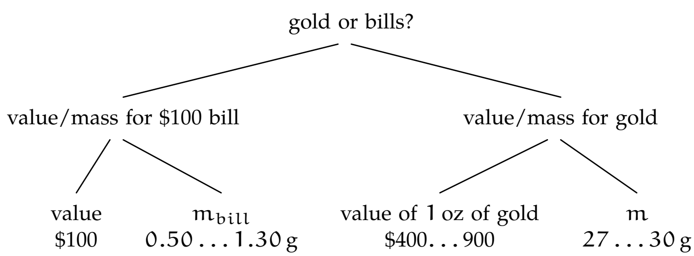

The next estimate is the value per mass of gold. I can be as accurate as I want in converting from ounces to grams. But I'll be lazy and try to remember the value while including uncertainty to reflect the fallibility of memory; let's say that $1oz = 27\ldots30\,g$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400…900. The mass and value ranges combine to give \$14…32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

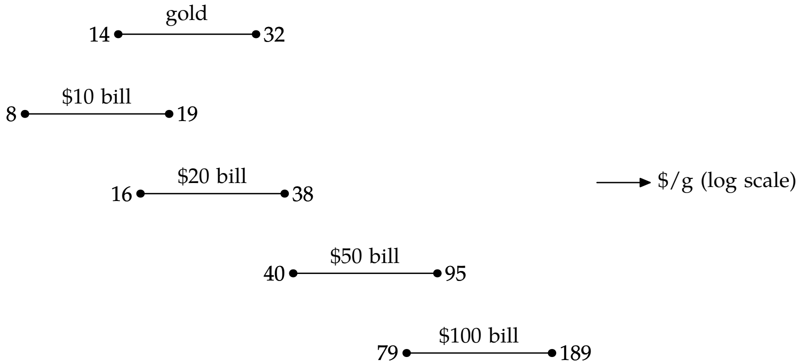

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

8.2.3 Oil imports using plausible ranges

To confirm these lessons, examine the benefit of divide-and-conquer reasoning in the example from Section 1.4: estimating the annual US oil imports. To quantify the benefit, I compare my plausible ranges before and after using divide and conquer.

Before I use divide and conquer, I have almost no idea what the oil imports are, and I am scared even to guess. To nudge me along, I imagine a mugger demanding, 'Your guess or your life!' In which case I counteroffer with, 'Can I give you a range instead of a number? I'd be surprised if the annual imports are less than10^7 barrels/yr or more than 10^{12} barrels/yr.' The imaginary mugger, being my own creation, always accepts my offer.

> **Problem 8.3 Your range**
> What is your plausible range for the annual oil imports?

I need little prodding to narrow my plausible range using divide-and-conquer reasoning. It required making several estimates:

1. N_{people}: US population;
2. f_{car}: cars per person;
3. l: average distance that a car is driven
4. m: average gas mileage;
5. V: volume of a barrel;
6. f_{other}: factor to multiply auto consumption to include all other consumption;
7. f_{imported}: fraction of oil that is imported.

> **Problem 8.4 Your ranges**
> Give your plausible range for each quantity, i.e. the range for which you assign a two-thirds probability that the true value lies within the range.

Here are my plausible ranges with a few notes of explanation:

1. N_{people}: 290–310 million. I recently read in the newspaper that the US population just reached the milestone of 300 million. How much should I believe what I read in the paper? The media lie when it serves the powerful, but I cannot find any reason to lie about the US

population, so I trust the figure, and throw in a bit of uncertainty to reflect the difficulties encountered in counting the population (e.g. what about undocumented immigrants, who are unlikely to fill out census forms?).

2. f_{car}: 0.5–1.5.
3. l: $7 \cdot 10^3$–$20 \cdot 10^3$ mi. Some books assessing used cars consider a low-mileage car to have less than 10^4 mi per year of age. So I guess that the average is somewhat larger than 10^4 mi/yr. But I am not confident of my recollection or the deduction, so my plausible range spans a factor of 3.
4. m: 15–40 miles/gallon;
5. V: 30–60 gallons;
6. f_{other}: 1.5–3;
7. f_{imported}: 0.3–0.8.

▶ *What is the resulting plausible range for the oil imports?*

Now combine the ranges using the method we used for the area of a sheet of A4 paper. That method produces the following plausible range:

$$1.0 \ldots 3.1 \ldots 9.6 \cdot 10^9 \text{ barrels/year.}$$

Compare this range to the range for the off-the-cuff guess $10^7 \ldots 10^{12}$ barrels/yr. That range spanned a factor of 10^5 whereas the improved range spans a mere factor of 10 – thanks to divide-and-conquer reasoning.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

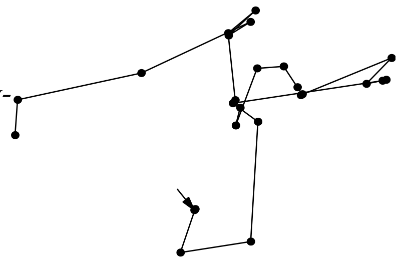

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

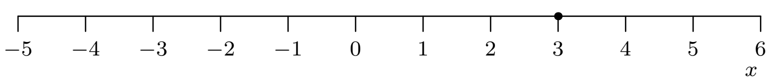

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

The result that $\langle x^2 \rangle$ is proportional to time applies not only to the one-dimensional random walk. Here's an example in two dimensions. Suppose that the particle's position is $(5, 2)$, so $\langle x^2 \rangle = 29$. After one step, it has four equally likely positions:

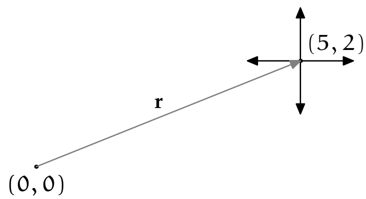

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6, 2)$ and $(4, 2)$. The expected squared distance is

$$\langle x^2 \rangle = \frac{1}{2}(40 + 20) = 30,$$

which is one more than the previous value of $\langle x^2 \rangle$. Since nothing is special about horizontal motion compared to vertical motion – symmetry! – the same result holds for vertical motion. So, averaging over the four possible locations produces an expected squared distance of 30.

For two dimensions, the pattern is:

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

No step in the analysis depended on being in only two dimensions. In fancy words, the derivation and the result are invariant to change of dimensionality. In plain English, this result also works in three dimensions.

In a standard walk in a straight line, $\langle x \rangle \propto$ time. Note the single power of x. The interesting quantity in a regular walk is not x itself, since it can grow without limit and is not invariant, but the ratio x/t, which is invariant to changes in t. This invariant is known as the speed.

In a random walk, where $\langle x^2 \rangle \propto t$, the interesting quantity is $\langle x^2 \rangle / t$. The expected squared position is not invariant to changes in t. However, the ratio $\langle x^2 \rangle / t$ is invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* and is often denoted D. It has dimensions of L^2T^{-1}.

This qualitative difference between a random and a regular walk makes intuitive sense. A random walker, for example a gas molecule or a very drunk person, moves back and forth, sometimes making progress in one direction, and other times undoing that progress. So, in order to travel the same distance, a random walker should require longer than a regular walker requires. The relation $\langle x^2 \rangle / t \sim D$ confirms and sharpens this intuition. The time for a random walker to travel a distance l is $t \sim l^2/D$, which grows quadratically rather than linearly with distance.

8.3.2 Diffusion equation

The preceding conclusion about random walks is sufficient to derive the diffusion equation, which describes how charge (electrons) move in a wire, how heat conducts through solid objects, and how gas molecules travel. Imagine then a gas of particles with each particle doing a random walk in one dimension. What is the equation that describes how the concentration, or number density, varies with time?

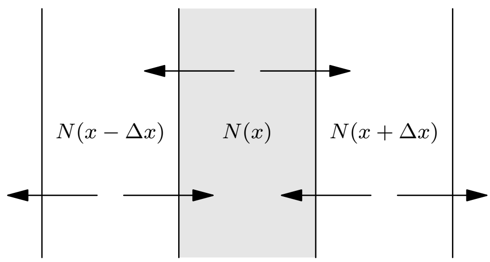

Divide the one-dimensional world into slices of width Δx, where Δx is the mean free path. Then look at the slices at $x - \Delta x$, x, and $x + \Delta x$. In every time step, one-half the molecules in each slice move left, and one-half move right. So the number of molecules in the x slice changes from $N(x)$ to

$$\frac{1}{2}(N(x-\Delta x) + N(x+\Delta x)).$$

The change in N is

$$\begin{aligned}\Delta N &= \frac{1}{2}(N(x-\Delta x) + N(x+\Delta x)) - N(x)\\ &= \frac{1}{2}(N(x-\Delta x) - 2N(x) + N(x+\Delta x)).\end{aligned}$$

This last relation can be rewritten as

$$\Delta N \sim (N(x+\Delta x) - N(x)) - (N(x) - N(x+\Delta x)).$$

In terms of derivatives, it is

$$\Delta N \sim (\Delta x)^2 \frac{\partial^2 N}{\partial x^2}.$$

The slices are separated by a distance such that most of the molecules travel from one piece to the neighboring piece in the time step τ. If τ is the time between collisions – the mean free time – then the distance is the mean free path λ. Thus

$$\frac{\Delta N}{\tau} \sim \frac{\lambda^2}{\tau}\frac{\partial^2 N}{\partial x^2},$$

or

$$\dot{N} \sim D\frac{\partial^2 N}{\partial x^2}$$

where $D \sim \lambda^2/\tau$ is a diffusion constant.

This partial-differential equation has interesting properties. The second spatial derivative means that a linear spatial concentration gradient remains unchanged. Its second derivative is zero so its time derivative must be zero. Diffusion fights curvature – roughly speaking, the second derivative – and does not try to change the gradient directly.

8.3.3 Keeping warm

One consequence of the diffusion equation is an analysis of how to keep warm on a cold day. To quantify keeping warm, or feeling cold, we need to calculate the heat flux: the energy flowing per unit area per unit time. Start with the definition of flux. Flux (of anything) is defined as

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{area} \times \text{time}}.$$

The flux depends on the density of stuff and on how fast the stuff travels:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass per temperature, $c_\mathrm{p}T$ is the thermal energy per mass, and $\rho c_\mathrm{p}T$ is therefore the thermal energy per volume. The speed is the 'speed' of diffusion. To diffuse a distance l takes time $t \sim l^2/D$, making the speed l/t or D/l. The l in the denominator indicates that, as expected, diffusion is slow over long distances. For heat diffusion, the diffusion constant is denoted κ and called the thermal diffusivity. So the speed is l/κ.

Combine the thermal energy per volume with the diffusion speed:

$$\text{thermal flux} = \rho c_{\mathrm{p}} T \times \frac{\kappa}{l}.$$

The product $\rho c_{\mathrm{p}}\kappa$ occurs so frequently that it is given a name: the thermal conductivity K. The ratio T/l is a lumped version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

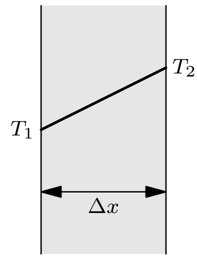

$$F = K\frac{\Delta T}{\Delta x}.$$

With one side held at T_1 and the other at T_2, the temperature gradient is $(T_2 - T_1)/\Delta x$.

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_{\mathrm{p}}\kappa$. To do so, put all the pieces together:

$$\begin{aligned} \rho &\sim 1\,\mathrm{kg\,m^{-3}}, \\ c_{\mathrm{p}} &\sim 10^3\,\mathrm{J\,kg^{-1}\,K^{-1}}, \\ \kappa &\sim 1.5\cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}, \end{aligned}$$

where we are guessing that $\kappa = \nu$ (because both κ and ν are diffusion constants). Then

$$K = \rho c_{\mathrm{p}}\kappa \sim 0.02\,\mathrm{W\,m^{-1}\,K^{-1}}.$$

Using this value we can estimate the heat loss on a cold day. Let's say that your skin is at $T_2 = 30\,^\circ\mathrm{C}$ and the air outside is $T_1 = 0\,^\circ\mathrm{C}$, making $\Delta T = 30\,\mathrm{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\mathrm{W\,m^{-1}\,K^{-1}} \times \frac{30\,\mathrm{K}}{2\cdot 10^{-3}\,\mathrm{m}} \sim 300\,\mathrm{W\,m^{-2}}.$$

Damn, we want a power rather than a power per area. Ah, flux is power per area, so just multiply by a person's surface area: roughly 2 m tall and 0.5 m wide, with a front and a back. So the surface area is about $2\,\mathrm{m}^2$. Thus, the power lost is

$$P \sim FA = 300\,\mathrm{W\,m^{-2}} \times 2\,\mathrm{m}^2 = 600\,\mathrm{W}.$$

No wonder a winter day wearing only thin pants and shirt feels so cold: 600 W is large compared to human power levels. Sitting around, a person produces 100 W of heat (the basal metabolic rate). When 600 W escapes, one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\mathrm{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

9
Springs

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

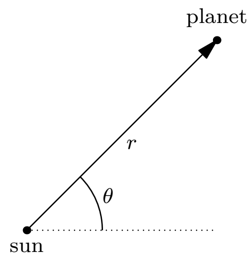

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law, angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2 \left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

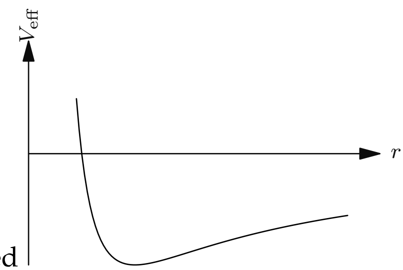

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

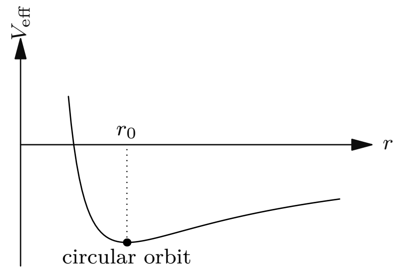

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

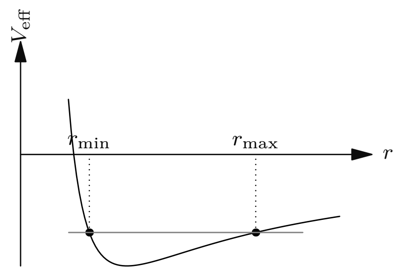

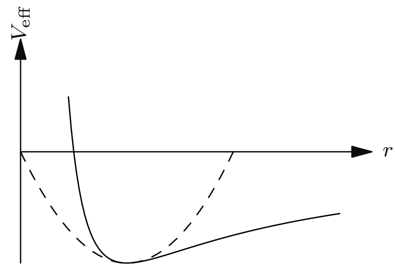

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r-r_0)^2\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r-r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

$$V = \frac{1}{2}\frac{k}{m}(x-a)^2, \tag{9.14}$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$\begin{aligned} x &\leftrightarrow r. \\ a &\leftrightarrow r_0. \\ k/m &\leftrightarrow GM/r_0^3. \end{aligned} \tag{9.15}$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \tag{9.16}$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

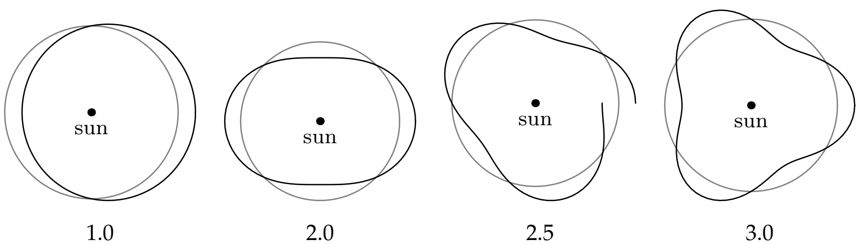

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \tag{9.17}$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

$$\frac{v^2}{r_0} = \frac{GM}{r_0^2}. \tag{9.18}$$

To manufacture v/r_0 on the left side, divide both sides by r_0 and take the square root:

$$\underbrace{\frac{v}{r_0}}_{\omega_{\text{orbit}}} = \sqrt{\frac{GM}{r_0^3}}. \tag{9.19}$$

This expression is also ω_{perturb}. Therefore, $\omega_{\text{orbit}} = \omega_{\text{perturb}}$.

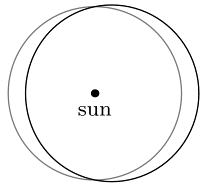

This result explains the elliptical (Kepler) orbits. First, the ratio is an integer, so the orbit is closed (compare the orbit with the ratio 2.5) – as it should be. Second, the ratio is 1, so the orbit's center is slightly away from the sun – as it should for a planetary orbit (the sun is at one focus of the ellipse, not at the center). The surprising conclusion of all the analysis is that a planetary orbit contains a spring; once this fact is appreciated, a spring analysis allows us to understand the complicated orbital motion without solving complicated differential equations.

As a bonus, the effective potential has the same form as the energy in hydrogen. Therefore, that energy also looks like a parabola near its minimum. The consequence is that a chemical bond acts like a spring (for small extensions). This second application is not a mere coincidence. Near a minimum, almost every function looks a parabola, so almost every physical system contains a spring. Springs are everywhere!

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

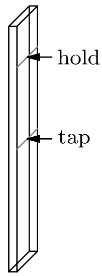

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

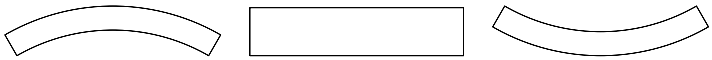

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

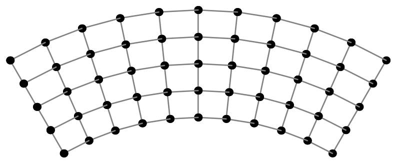

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

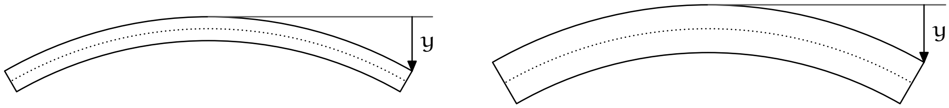

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

In general, $m \propto h$ so

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{h^3}{h}} = h. \tag{9.27}$$

Frequency is proportional to thickness!

9.2.2 Xylophone

Let's check this analysis by looking at its consequences and comparing with experimental data from a home experiment. My daughter got a toy xylophone from her uncle. Its slats have these dimensions:

	ℓ
C	12.2 cm
D	11.5
E	10.9
F	10.6
G	10.0
A	9.4
B	8.9
C′	8.6

Our analysis of how frequency depends on thickness can explain this pattern of how frequency depends on length. The method is to use dimensional analysis with proportional reasoning (scaling).

Rather than finding the frequency direction, I analyze the stiffness. The mass is easy, so split that part off of the calculation of the frequency. The block's spring constant k depends on its material properties – here, the Young's modulus Y – and on its dimensions. So the variables are k, Y, and length l, width w, and thickness (height) h.

▶ *How many independent dimensions are contained in those variables? How many independent dimensionless groups can be formed from those variables?*

These five variables are composed of two independent dimensions. These dimensions could be length and force: Stiffness is force per length, and Young's modulus is force per area. Five variables based on two independent dimensions form three independent dimensionless groups. The goal is to find k, so I include k in only one group. That group contains Y to divide out the dimensions of mass. Since Young's modulus is force per area, and stiffness is force per length, the ratio k/Yh is dimensionless. The three lengths for the size of the block easily make two more dimensionless groups: for example, h/l and w/l. Then

$$\frac{k}{Yh} = f\left(\frac{h}{l}, \frac{w}{l}\right). \tag{9.28}$$

▶ *Guess the function f (except for a dimensionless constant).*

What we know about stiffness versus thickness, along with proportional reasoning, is enough to solve for f, except for a dimensionless constant. Proportional reasoning helps determine the dependence on the dimensionless group w/l. Imagine doubling the width w. Equivalently, I glue together two identical blocks along the long, thin edge. When the new block is bent, the individual blocks contains equal energy, so the new block contains twice the energy of an original block. Therefore, doubling the width doubles the stiffness; in symbols, $k \propto w$. In the general form

$$\frac{k}{Yh} = f\left(\frac{h}{l}, \frac{w}{l}\right). \tag{9.29}$$

w appears only in the group w/l, and k appears on the right in the first power. So the general form simplifies to

$$\frac{k}{Yh} = \frac{w}{l} \cdot f\left(\frac{h}{l}\right). \tag{9.30}$$

To guess the new function f, I use what I know about stiffness versus thickness, that $k \propto h^3$. Therefore the left side, k/Yh, is proportional to h^2. On the right side the only source of h is from f, which can play with h but only via the ratio h/l. So

$$f\left(\frac{h}{l}\right) \sim \left(\frac{h}{l}\right)^2. \tag{9.31}$$

Combining these deductions gives

$$\frac{k}{Yh} \sim \frac{w}{l}\left(\frac{h}{l}\right)^2 = \frac{wh^2}{l^3} \tag{9.32}$$

and

$$k \sim Yw\left(\frac{h}{l}\right)^3. \tag{9.33}$$

The stiffness and mass determine the frequency. The mass is $m = \rho wlh$. So

$$\omega \sim \sqrt{\frac{k}{m}} \sim \sqrt{\frac{Y}{\rho}}\,\frac{h}{l^2}. \tag{9.34}$$

As a quick check, this result is consistent with the earlier calculation that frequency is proportional to thickness. And it contains a new result: $\omega \propto l^{-2}$.

Problem 9.1 Effect of width
Is it physically plausible that the width w does not affect the frequency ω?

▶ *Is this data consistent with the prediction that $\omega \propto l^{-2}$?*

Before doing an extensive analysis, I check the easy case of the octave. The lower and higher C notes are a factor of 2 apart in frequency. If the scaling prediction is correct, the respective slat lengths should be a factor of $\sqrt{2}$ apart. The length ratio is $12.2/8.6 \sim 1.419$, which is very close to $\sqrt{2}$. The general pattern is that fl^2 should be invariant. To check, here is the same table with frequencies, which are computed by assuming that the A above middle C is at 440 Hz (concert A), and with a column for fl^2:

	ℓ	f
C	12.2	261.6
D	11.5	293.6
E	10.9	329.6
F	10.6	349.2
G	10.0	392.0
A	9.4	440.0
B	8.9	493.8
C′	8.6	523.2

The proposed invariant is, experimentally, almost constant.

9.3 Waves

Ocean covers most of the earth, and waves roam most of the ocean. Waves also cross puddles and ponds. What makes them move, and what determines their speed? By applying and extending the techniques of approximation, we analyze waves. For concreteness, this section refers mostly to water waves but the results apply to any fluid.

9.3.1 Dispersion relations

The most organized way to study waves is to use dispersion relations. A dispersion relation states what values of frequency and wavelength a wave can have. Instead of the wavelength λ, dispersion relations usually connect frequency ω, and wavenumber k, which is defined as $2\pi/\lambda$. This preference has an basis in order-of-magnitude reasoning. Wavelength is the the distance the wave travels in a full period, which is 2π radians of oscillation. Although 2π is dimensionless, it is not the ideal dimensionless number, which is unity. In 1 radian of oscillation, the wave travels a distance

$$\bar{\lambda} \equiv \frac{\lambda}{2\pi}. \tag{9.35}$$

The bar notation, meaning 'divide by 2π', is chosen by analogy with h and $\hbar$. The one-radian forms such as $\hbar$ are more useful for approximations than the 2π-radian forms. The Bohr radius, in a form where the dimensionless constant is unity, contains $\hbar$ rather than h. Most results with waves are similarly simpler using $\bar{\lambda}$ rather than λ. A further refinement is to use its inverse, the wavenumber $k = 1/\bar{\lambda}$. This choice, which has dimensions of inverse length, parallels the definition of angular frequency ω, which has dimensions of inverse time. A relation that connects ω and k is likely to be simpler than one connecting ω and $\bar{\lambda}$, although either is simpler than one connecting ω and λ.

The simplest dispersion relation describes electromagnetic waves in a vacuum. Their frequency and wavenumber are related by the dispersion relation

$$\omega = ck, \tag{9.36}$$

which states that waves travel at velocity $\omega/k = c$, independent of frequency. Dispersion relations contain a vast amount of information about waves. They contain, for example, how fast crests and troughs travel: the phase velocity. They contain how fast wave packets travel: the group velocity. They contain how these velocities depend on frequency: the dispersion. And they contain the rate of energy loss: the attenuation.

9.3.2 Phase and group velocities

The usual question with a wave is how fast it travels. This question has two answers, the phase velocity and the group velocity, and both depend on the dispersion relation. To get a feel for how to use dispersion relations (most of the chapter is about how to calculate them), we discuss the simplest examples that illustrate these two velocities. These analyses start with the general form of a traveling wave:

$$f(x,t) = \cos(kx - \omega t), \tag{9.37}$$

where f is its amplitude.

Phase velocity is an easier idea than group velocity so, as an example of

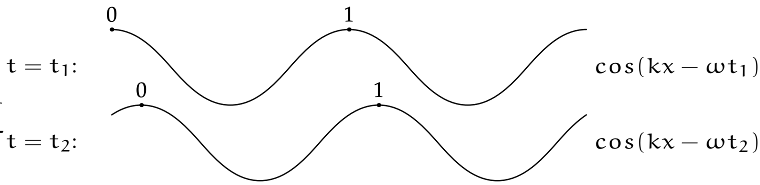

divide-and-conquer reasoning and of maximal laziness, study it first. The phase, which is the argument of the cosine, is $kx-\omega t$. A crest occurs when the phase is zero. In the top wave, a crest occurs when $x = \omega t_1/k$. In the bottom wave, a crest occurs when $x = \omega t_2/k$. The difference

$$\frac{\omega}{k}(t_2 - t_1) \tag{9.38}$$

is the distance that the crest moved in time $t_2 - t_1$. So the phase velocity, which is the velocity of the crests, is

$$v_{\text{ph}} = \frac{\text{distance crest shifted}}{\text{time taken}} = \frac{\omega}{k}. \tag{9.39}$$

To check this result, check its dimensions: ω provides inverse time and $1/k$ provides length, so ω/k is a speed.

Group velocity is trickier. The word 'group' suggests that the concept involves more than one wave. Because two is the first whole number larger than one, the simplest illustration uses two waves. Instead of being a function relating ω and k, the dispersion relation here is a list of allowed (k, ω) pairs. But that form is just a discrete approximation to an official dispersion relation, complicated enough to illustrate group velocity and simple enough to not create a forest of mathematics. So here are two waves with almost the same wavenumber and frequency:

$$\begin{aligned} f_1 &= \cos(kx - \omega t), \\ f_2 &= \cos((k + \Delta k)x - (\omega + \Delta\omega)t), \end{aligned} \tag{9.40}$$

where Δk and $\Delta\omega$ are small changes in wavenumber and frequency, respectively. Each wave has phase velocity $v_{\text{ph}} = \omega/k$, as long as Δk and $\Delta\omega$ are tiny. The figure shows their sum.

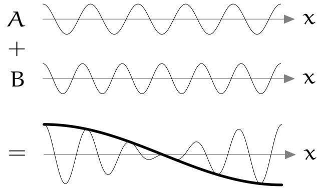

The point of the figure is that two cosines with almost the same spatial frequency add to produce an envelope (thick line). The envelope itself looks like a cosine. After waiting a while, each wave changes because of the ωt or $(\omega + \Delta\omega)t$ terms in their phases. So the sum and its envelope change to this:

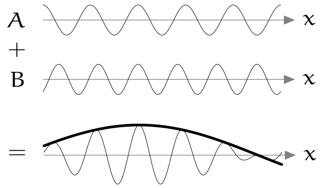

The envelope moves, like the crests and troughs of any wave. It is a wave, so it has a phase velocity, which motivates the following definition:

Group velocity is the phase velocity of the envelope. (9.41)

With this pictorial definition, you can compute group velocity for the wave $f_1 + f_2$. As suggested in the figures, adding two cosines produces a a slowly varying envelope times a rapidly oscillating inner function. This trigonometric identity gives the details:

$$\cos(A+B) = \underbrace{2\cos\left(\frac{B-A}{2}\right)}_{\textit{envelope}} \times \underbrace{\cos\left(\frac{A+B}{2}\right)}_{\textit{inner}}. \tag{9.42}$$

If $A \approx B$, then $A - B \approx 0$, which makes the envelope vary slowly. Apply the identity to the sum:

$$f_1 + f_2 = \cos(\underbrace{kx - \omega t}_{A}) + \cos(\underbrace{(k+\Delta k)x - (\omega + \Delta\omega)t}_{B}). \tag{9.43}$$

Then the envelope contains

$$\cos\left(\frac{B-A}{2}\right) = \cos\left(\frac{x\Delta k - t\Delta\omega}{2}\right). \tag{9.44}$$

The envelope represents a wave with phase

$$\frac{\Delta k}{2}x - \frac{\Delta\omega}{2}t. \tag{9.45}$$

So it is a wave with wavenumber $\Delta k/2$ and frequency $\Delta\omega/2$. The envelope's phase velocity is the group velocity of $f_1 + f_2$:

$$v_g = \frac{\text{frequency}}{\text{wavenumber}} = \frac{\Delta\omega/2}{\Delta k/2} = \frac{\Delta\omega}{\Delta k}. \tag{9.46}$$

In the limit where $\Delta k \to 0$ and $\Delta\omega \to 0$, the group velocity is

$$v_g = \frac{\partial\omega}{\partial k}. \tag{9.47}$$

It is usually different from the phase velocity. A typical dispersion relation, which appears several times in this chapter, is $\omega \propto k^n$. Then $v_{ph} = \omega/k = k^{n-1}$ and $v_g \propto nk^{n-1}$. So their ratio is

$$\frac{v_g}{v_{ph}} = n. \qquad \text{(for a power-law relation)} \tag{9.48}$$

Only when $n = 1$ are the two velocities equal. Now that we can find wave velocities from dispersion relations, we return to the problem of finding the dispersion relations.

9.3.3 Dimensional analysis

A dispersion relation usually emerges from solving a wave equation, which is an unpleasant partial differential equation. For water waves, a wave equation emerges after linearizing the equations of hydrodynamics and neglecting viscosity. This procedure is mathematically involved, particularly in handling the boundary conditions. Alternatively, you can derive dispersion relations using dimensional analysis, then complete and complement the derivation with physical arguments. Such methods usually cannot evaluate the dimensionless constants, but the beauty of studying waves is that, as in most problems involving springs and oscillations, *most of these constants are unity.*

How do frequency and wavenumber connect? They have dimensions of T^{-1} and L^{-1}, respectively, and cannot form a dimensionless group without help. So include more variables. What physical properties of the system determine wave behavior? Waves on the open ocean behave differently from waves in a bathtub, perhaps because of the difference in the depth of water h. The width of the tub or ocean could matter, but then the problem becomes two-dimensional wave motion. In a first analysis, avoid that complication and consider waves that move in only one dimension, perpendicular to the width of the container. Then the width does not matter.

To determine what other variables are important, use the principle that waves are like springs, because *every physical process contains a spring*. This blanket statement cannot be strictly correct. However, it is useful as a broad generalization. To get a more precise idea of when this assumption is useful, consider the characteristics of spring motion. First, springs have an equilibrium position. If a system has an undisturbed, resting state, consider looking for a spring. For example, for waves on the ocean, the undisturbed state is a calm, flat ocean. For electromagnetic waves – springs are not confined to mechanical systems – the resting state is an empty vacuum with no radiation. Second, springs oscillate. In mechanical systems, oscillation depends on inertia to carry the mass beyond the equilibrium position. Equivalently, it depends on kinetic energy turning into potential energy, and vice versa. Water waves store potential energy in the disturbance of the surface and kinetic energy in the motion of the water. Electromagnetic waves store energy in the electric and magnetic fields. A magnetic field is generated by moving or spinning charges, so the magnetic field is a reservoir of kinetic (motion) energy. An electric field is generated by stationary charges and has an associated potential, so the electric field is the reservoir of potential energy. With these identifications, the electromagnetic field acts like a set of springs, one for each radiation frequency. These examples are positive examples. A negative example – a marble oozing its way through glycerin – illustrates that springs are not always present. The marble moves so slowly that the kinetic energy of the corn syrup, and therefore its inertia, is miniscule and irrelevant. This system has no reservoir of kinetic energy, for the kinetic energy is merely dissipated, so it does not contain a spring.

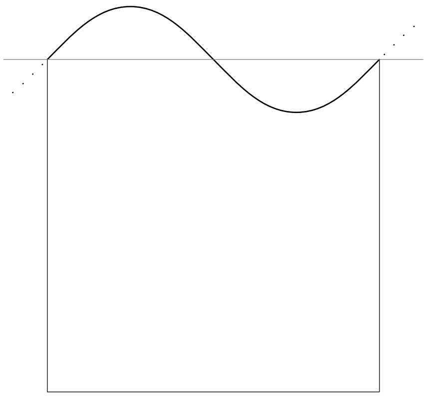

Waves have the necessary reservoirs to act like springs. The surface of water is flat in its lowest-energy state. Deviations, also known as waves, are opposed by a restoring force. Distorting the surface is like stretching a rubber sheet: Surface tension resists the distortion. Distorting the surface also requires raising the average water level, a change that gravity resists.

The average *height* of the surface does not change, but the average depth of the water does. The higher column, under the crest, has more water than the lower column, under the trough. So in averaging to find the average depth, the higher column gets a slightly higher weighting. Thus the average depth increases. This result is consistent with experience: It takes energy to make waves.

The total restoring force includes gravity and surface tension so, in the list of variables, include surface tension (γ) and gravity (g).

In a wave, like in a spring, the restoring force fights inertia, represented here by the fluid density. The gravitational piece of the restoring force does not care about density: Gravity's stronger pull on denser substances is exactly balanced by their greater inertia. This exact cancellation is a restatement of the equivalence principle, on which Einstein based the theory of general relativity [4, 5]. In pendulum motion, the mass of the bob drops out of the final solution for the same reason. The surface-tension piece of the restoring force, however, does not change when density changes. The surface tension itself, the fluid property γ, depends on density because it depends on the spacing of atoms at the surface. That dependence affects γ. However, once you know γ you can compute surface-tension forces without knowing the density. Since ρ does not affect the surface-tension force but affects the inertia, it affects the properties of waves in which surface tension provides a restoring force. Therefore, include ρ in the list.

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

To simplify the analysis, assume that the waves do not lose energy. This choice excludes viscosity from the set of variables. To further simplify, exclude the speed of sound. This approximation means ignoring sound waves, and is valid as long as the flow speeds are slow compared to the speed of sound. The resulting ratio,

$$\mathcal{M} \equiv \frac{\text{flow speed}}{\text{sound speed}}, \tag{9.49}$$

is dimensionless and, because of its importance, is given a name: the Mach number. Finally, assume that the wave amplitude ξ is small compared to its wavelength and to the depth of the container. The table shows the list of variables. Even with all these restrictions, which significantly simplify the analysis, the results explain many phenomena in the world.

These six variables built from three fundamental dimensions produce three dimensionless groups. One group is easy: the wavenumber k is an inverse length and the depth h is a length, so

$$\Pi_1 \equiv kh. \tag{9.50}$$

This group is the dimensionless depth of the water: $\Pi_1 \ll 1$ means shallow and $\Pi_1 \gg 1$ means deep water. A second dimensionless group comes from gravity. Gravity, represented by g, has the same dimensions as ω^2, except for a factor of length. Dividing by wavenumber fixes this deficit:

$$\Pi_2 = \frac{\omega^2}{gk}. \tag{9.51}$$

Without surface tension, Π_1 and Π_2 are the only dimensionless groups, and neither group contains density. This mathematical result has a physical basis. Without surface tension, the waves propagate because of gravity alone. The equivalence principle says that gravity affects motion independently of density. Therefore, density cannot – and does not – appear in either group.

Now let surface tension back into the playpen of dimensionless groups. It must belong in the third (and final) group Π_3. Even knowing that γ belongs to Π_3 still leaves great freedom in choosing its form. The usual pattern is to find the group and then interpret it, as we did for Π_1 and

Π_2. Another option is to begin with a physical interpretation and use the interpretation to construct the group. Here you can construct Π_3 to measure the relative importance of surface-tension and gravitational forces. Surface tension γ has dimensions of force per length, so producing a force requires multiplying by a length. The problem already has two lengths: wavelength (represented via k) and depth. Which one should you use? The wavelength probably always affects surface-tension forces, because it determines the curvature of the surface. The depth, however, affects surface-tension forces only when it becomes comparable to or smaller than the wavelength, if even then. You can use both lengths to make γ into a force: for example, $F \sim \gamma\sqrt{h/k}$. But the analysis is easier if you use only one, in which case the wavelength is the preferable choice. So $F_\gamma \sim \gamma/k$. Gravitational force, also known as weight, is $\rho g \times$ volume. Following the precedent of using only k to produce a length, the gravitational force is $F_\mathrm{g} \sim \rho g/k^3$. The dimensionless group is then the ratio of surface-tension to gravitational forces:

$$\Pi_3 \equiv \frac{F_\gamma}{F_\mathrm{g}} = \frac{\gamma/k}{\rho g/k^3} = \frac{\gamma k^2}{\rho g}. \tag{9.52}$$

This choice has, by construction, a useful physical interpretation, but many other choices are possible. You can build a third group without using gravity: for example, $\Pi_3 \equiv \gamma k^3/\rho\omega^2$. With this choice, ω appears in two groups: Π_2 and Π_3. So it will be hard to solve for it. The choice made for P_3, besides being physically useful, quarantines ω in one group: a useful choice since ω is the goal.

Now use the groups to solve for frequency ω as a function of wavenumber k. You can instead solve for k as a function of ω, but the formulas for phase and group velocity are simpler with ω as a function of k. Only the group Π_2 contains ω, so the general dimensionless solution is

$$\Pi_2 = f(\Pi_1, \Pi_3), \tag{9.53}$$

or

$$\frac{\omega^2}{gk} = f\left(kh, \frac{\gamma k^2}{\rho g}\right). \tag{9.54}$$

Then

$$\omega^2 = gk \cdot f(kh, \frac{\gamma k^2}{\rho g}). \tag{9.55}$$

This relation is valid for waves in shallow or deep water (small or large Π_1); for waves propagated by gravity or by surface tension (small or large Π_3); and for waves in between.

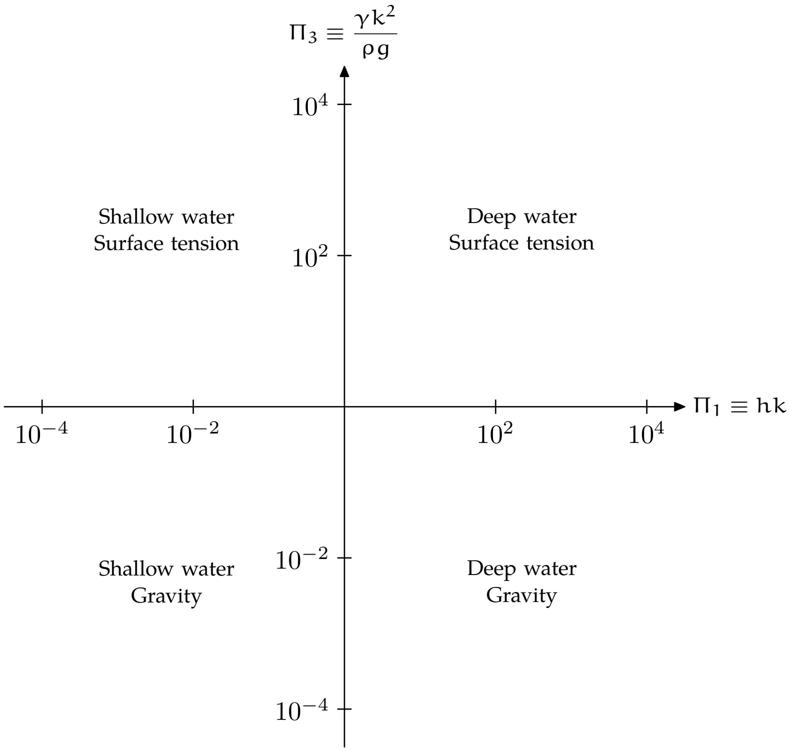

The figure shows how the two groups Π_1 and Π_3 divide the world of waves into four regions. We study each region in turn, and combine the analyses to understand the whole world (of waves). The group Π_1 measures the depth of the water: Are the waves moving on a puddle or an ocean? The group Π_3 measures the relative contribution of surface tension and gravity: Are the waves ripples or gravity waves?

The division into deep and shallow water (left and right sides) follows from the interpretation of $\Pi_1 = kh$ as dimensionless depth. The division into surface-tension- and gravity-dominated waves (top and bottom halves) is more subtle, but is a result of how Π_3 was constructed. As a check, look at Π_3. Large g or small k result in the same consequence: small Π_3. Therefore the physical consequence of longer wavelength (smaller k) is similar to that of stronger gravity. So longer-wavelength waves are gravity waves. The large-Π_3 portion of the world (top half) is therefore labeled with surface tension.

The next figure shows how wavelength and depth (instead of the dimensionless groups) partition the world, and plots examples of different types of waves.

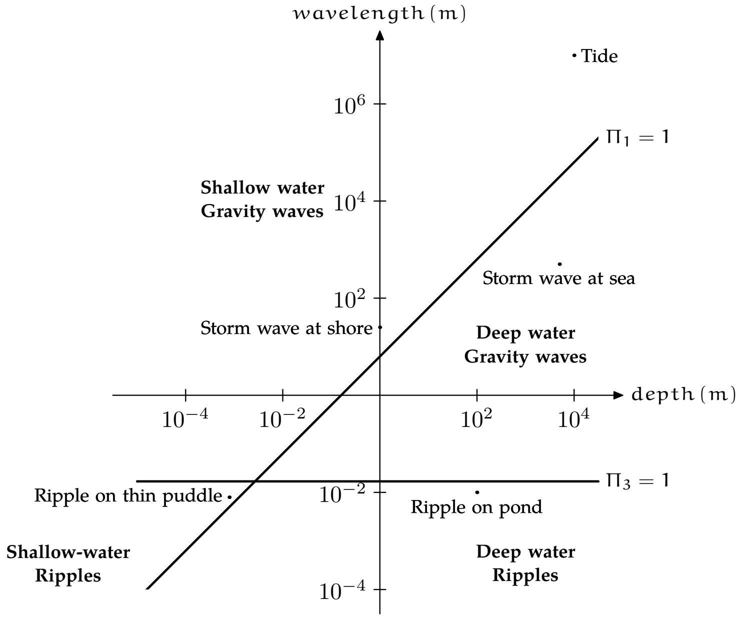

The thick dividing lines are based on the dimensionless groups $\Pi_1 = hk$ and $\Pi_3 = \gamma k^2/\rho g$. Each region contains one or two examples of its kind of waves. With $g = 1000\,\mathrm{cm\,s^{-1}}$ and $\rho \sim 1\,\mathrm{g\,cm^{-3}}$, the border wavelength between ripples and gravity waves is just over $\lambda \sim 1\,\mathrm{cm}$ (the horizontal, $\Pi_3 = 1$ dividing line).

The magic function f is still unknown to us. To determine its form and to understand its consequences, study f in limiting cases. Like a jigsaw-puzzle-solver, study first the corners of the world, where the physics is simplest. Then connect the corner solutions to get solutions valid along an edge, where the physics is the almost as simple as in a corner. Finally, crawl inward to assemble the complicated, complete solution. This extended example illustrates divide-and-conquer reasoning, and using limiting cases to choose pieces into which you break the problem.

9.3.4 Deep water

First study deep water, where $kh \gg 1$, as shaded in the map. Deep water is defined as water sufficiently deep that waves cannot feel the bottom of the ocean. How deep do waves' feelers extend? The only length scale in the waves is the wavelength, $\lambda = 2\pi/k$. The feelers therefore extend to a depth $d \sim 1/k$ (as always, neglect constants, such as 2π). This educated guess has a justification in Laplace's equation, which is the spatial part of the wave equation. Suppose that the waves are periodic in the x direction, and z measures depth below the surface, as shown in this figure:

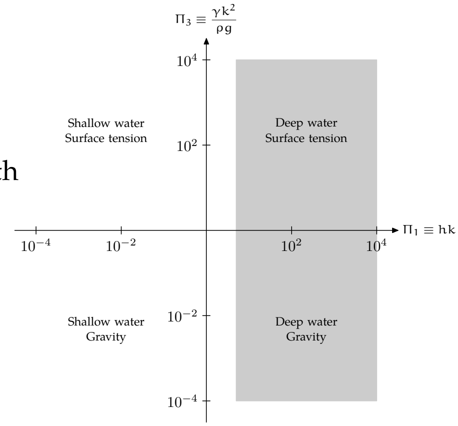

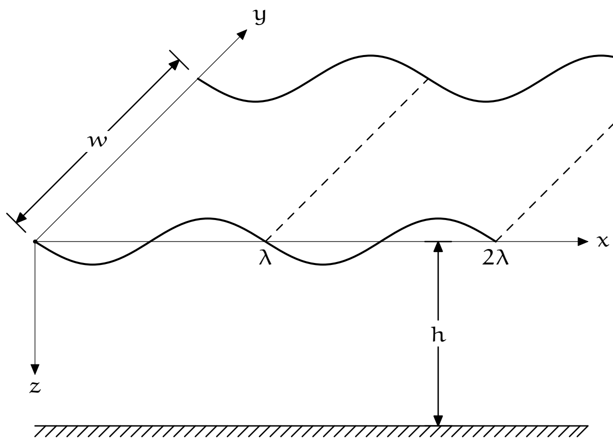

Then, Laplace's equation becomes

$$\frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial z^2} = 0, \tag{9.56}$$

where ϕ is the velocity potential. The $\partial^2\phi/\partial y^2$ term vanishes because nothing varies along the width (the y direction).

It's not important what exactly ϕ is. You can find out more about it in an excellent fluid-mechanics textbook, *Fluid Dynamics for Physicists* [6]; Lamb's *Hydrodynamics* [14]is a classic but difficult. For this argument, all that matters is that ϕ measures the effect of the wave and that ϕ satisfies Laplace's equation. The wave is periodic in the x direction, with a form such as $\sin kx$. Take

$$\phi \sim Z(z) \sin kx. \tag{9.57}$$

The function $Z(z)$ measures how the wave decays with depth.

The second derivative in x brings out two factors of k, *and a minus sign*:

$$\frac{\partial^2 \phi}{\partial x^2} = -k^2 \phi. \tag{9.58}$$

In order that this ϕ satisfy Laplace's equation, the z-derivative term must contribute $+k^2\phi$. Therefore,

$$\frac{\partial^2 \phi}{\partial z^2} = k^2 \phi, \tag{9.59}$$

so $Z(z) \sim e^{\pm kz}$. The physically possible solution – the one that does not blow up exponentially at the bottom of the ocean – is $Z(z) \sim e^{-kz}$. Therefore, relative to the effect of the wave at the surface, the effect of the wave at the bottom of the ocean is $\sim e^{-kh}$. When $kh \gg 1$, the bottom might as well be on the moon because it has no effect. The dimensionless factor kh – it must be dimensionless to sit alone in an exponent – compares water depth with feeler depth $d \sim 1/k$:

$$\frac{\text{water depth}}{\text{feeler depth}} \sim \frac{h}{1/k} = hk, \tag{9.60}$$

which is the dimensionless group Π_1.

In deep water, where the bottom is hidden from the waves, the water depth h does not affect their propagation, so h disappears from the list of relevant variables. When h goes, so does $\Pi_1 = kh$. There is one caveat. If Π_1 is the only group that contains k, then you cannot blithely discard Π_1 just because you no longer care about h. If you did, you would be discarding k and h, and make it impossible to find a dispersion relation (which connects ω and k). Fortunately, k appears in $\Pi_3 = \gamma k^2/\rho g$ as well as in Π_1. So in deep water it is safe to discard Π_1. This argument for the irrelevance of h is a physical argument. It has a mathematical equivalent in the language of dimensionless groups and functions. Because h has dimensions, the statement that 'h is large' is meaningless. The question is, 'large compared to what length?' With $1/k$ as the standard of comparison the meaningless 'h is large' statement becomes 'kh is large.' The product kh is the dimensionless group Π_1. Mathematically, you are assuming that the function $f(kh, \gamma k^2/\rho g)$ has a limit as $kh \to \infty$.

Without Π_1, the general dispersion relation simplifies to

$$\omega^2 = gk\, f_{\text{deep}}\left(\frac{\gamma k^2}{\rho g}\right). \tag{9.61}$$

This relation describes the deep-water edge of the world of waves. The edge has two corners, labeled by whether gravity or surface tension provides the restoring force. Although the form of f_{deep} is unknown, it is a simpler function than the original f, a function of two variables. To determine the form of f_{deep}, continue the process of dividing and conquering: Partition deep-water waves into its two limiting cases, gravity waves and ripples.

9.3.5 Gravity waves on deep water

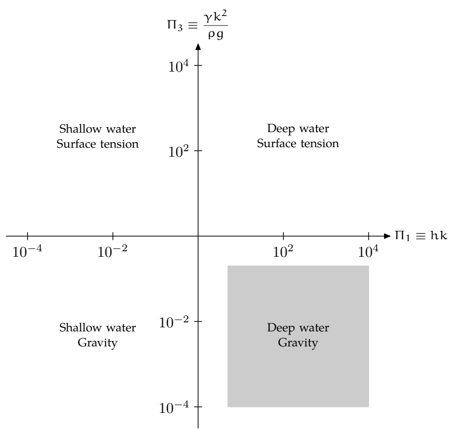

Of the two extremes, gravity waves are the more common. They include wakes generated by ships and most waves generated by wind. So specialize to the corner of the wave world where water is deep and gravity is strong. With gravity much stronger than surface tension, the dimensionless group $\Pi_3 = \gamma k^2/\rho g$ limits to 0. Since Π_3 is the product of several factors, you can achieve the limit in several ways:

1. Increase g (which is downstairs) by moving to Jupiter.
2. Reduce γ (which is upstairs) by dumping soap on the water.
3. Reduce k (which is upstairs) by studying waves with a huge wavelength.

In this limit, the general deep-water dispersion relation simplifies to

$$\omega^2 = f_{\text{deep}}(0)gk = C_1 gk, \tag{9.62}$$

where $f_{\text{deep}}(0)$ is an as-yet-unknown constant, C_1. The use of $f_{\text{deep}}(0)$ assumes that $f_{\text{deep}}(x)$ has a limit as $x \to 0$. The slab argument, which follows shortly, shows that it does. For now, in order to make progress, assume that it has a limit. The constant remains unknown to the lazy methods of dimensional analysis, because the methods sacrifice evaluation of dimensionless constants to gain comprehension of physics. Usually assume that such constants are unity. In this case, we get lucky: An honest calculation produces $C_1 = 1$ and

$$\omega^2 = 1 \times gk, \tag{9.63}$$

where the red $1\times$ indicates that it is obtained from honest physics.

Such results from dimensional analysis seem like rabbits jumping from a hat. The dispersion relation is correct, but your gut may grumble about this magical derivation and ask, 'But *why* is the result true?' A physical model of the forces or energies that drive the waves explains the origin of the dispersion relation. The first step is to understand the mechanism: How does gravity make the water level rise and fall? Taking a hint from the Watergate investigators,[1] we follow the water. The water in the crest does *not* move into the trough. Rather, the water in the crest, being higher, creates a pressure underneath it higher than that of the water in the trough, as shown in this figure:

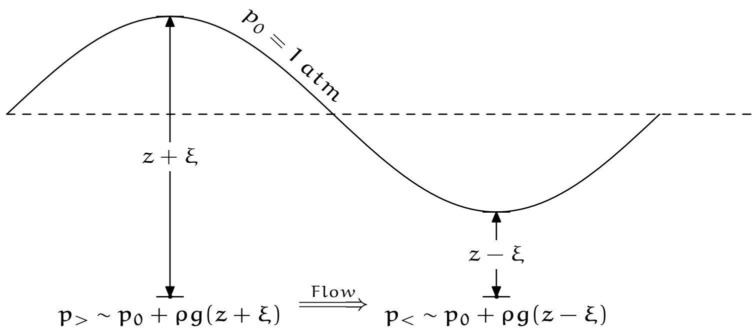

1 When the reporters Woodward and Bernstein [2]were investigating criminal coverups during the Nixon administration, they received help from the mysterious 'Deep Throat', whose valuable advice was to 'follow the money.'

The higher pressure forces water underneath the crest to flow toward the trough, making the water level there rise. Like a swing sliding past equilibrium, the surface overshoots the equilibrium level to produce a new crest and the cycle repeats.

The next step is to quantify the model by estimating sizes, forces, speeds, and energies. In Section 7.1 we analyzed a messy mortality curve by replacing it with a more tractable shape: a rectangle. The method of lumping worked there, so try it again. 'A method is a trick I use twice.' —George Polyà. Water just underneath the surface moves quickly because of the pressure gradient. Farther down, it moves more slowly. Deep down it does not move at all. Replace this smooth falloff with a step function: Pretend that water down to a certain depth moves as a block, while deeper water stays still:

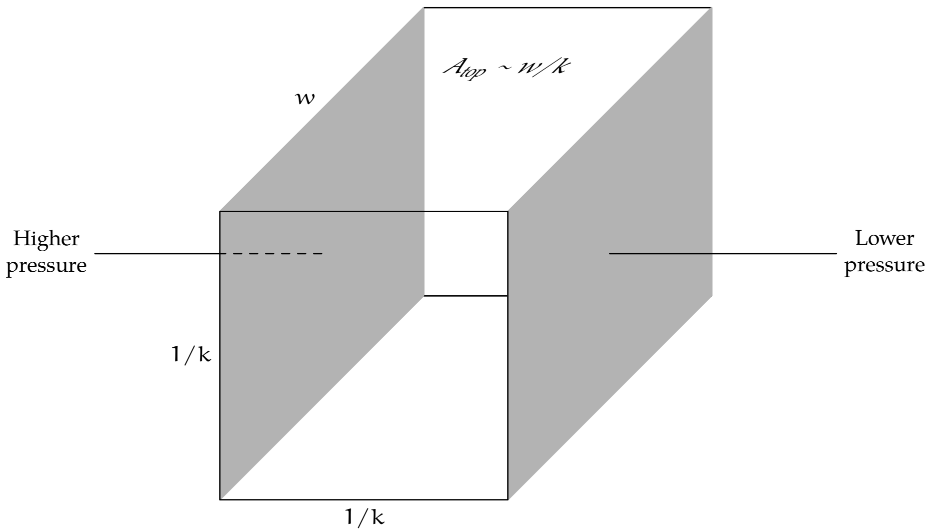

How deep should this slab of water extend? By the Laplace-equation argument, the pressure variation falls off exponentially with depth, with length scale $1/k$. So assume that the slab has a similar length scale, that it has depth $1/k$. What choice do you have? On an infinitely deep ocean, the only length scale is $1/k$. How long should the slab be? Its length should be roughly the peak-to-trough distance of the wave because the surface height changes significantly over that distance. This distance is $1/k$. Actually, it is π/k (one-half of a period), but ignore constants. All the constants combine into a giant constant at the end, which dimensional analysis cannot determine anyway, so discard it now! The slab's width w is arbitrary and cancels by the end of any analysis.

So the slab of water has depth $1/k$, length $1/k$, and width w. Estimate the forces acting on it by estimating the pressure gradients. Across the width of the slab (the y direction), the water surface is level, so the pressure is constant along the width. Into the depths (the z direction), the pressure varies because of gravity – the $\rho g h$ term from hydrostatics – but that variation is just sufficient to prevent the slab from sinking. We care about only the pressure difference across the length, the direction that the wave moves. This pressure difference depends on the height of the crest, ξ and is $\Delta p \sim \rho g \xi$. This pressure difference acts on a cross-section with area $A \sim w/k$ to produce a force

$$F \sim \underbrace{w/k}_{\textit{area}} \times \underbrace{\rho g \xi}_{\Delta p} = \rho g w \xi / k. \tag{9.64}$$

The slab has mass

$$m = \rho \times \underbrace{w/k^2}_{\textit{volume}}, \tag{9.65}$$

so the force produces an acceleration

$$a_{\text{slab}} \sim \underbrace{\frac{\rho g w \xi}{k}}_{\textit{force}} \bigg/ \underbrace{\frac{\rho w}{k^2}}_{\textit{mass}} = g \xi k. \tag{9.66}$$

The factor of g says that the gravity produces the acceleration. Full gravitational acceleration is reduced by the dimensionless factor ξk, which is roughly the slope of the waves.

The acceleration of the slab determines the acceleration of the surface. If the slab moves a distance x, it sweeps out a volume of water $V \sim xA$. This water moves under the trough, and forces the surface upward a distance V/A_{top}. Because $A_{\text{top}} \sim A$ (both are $\sim w/k$), the surface moves the same distance x that the slab moves. Therefore, the slab's acceleration a_{slab} equals the acceleration a of the surface:

$$a \sim a_{\text{slab}} \sim g \xi k. \tag{9.67}$$

This equivalence of slab and surface acceleration does not hold in shallow water, where the bottom at depth h cuts off the slab before $1/k$; that story is told in Section 9.3.12.

The slab argument is supposed to justify the deep-water dispersion relation derived by dimensional analysis. That relation contains frequency whereas acceleration relation does not. So massage it until ω appears. The acceleration relation contains a and ξ, whereas the dispersion relation does not. An alternative expression for the acceleration might make the acceleration relation more like the dispersion relation. With luck the expression will contain ω^2, thereby producing the hoped-for ω^2; as a bonus, it will contain ξ to cancel the ξ in the acceleration relation.

In simple harmonic motion (springs!), acceleration is $a \sim \omega^2\xi$, where ξ is the amplitude. In waves, which behave like springs, a is given by the same expression. Here's why. In time $\tau \sim 1/\omega$, the surface moves a distance $d \sim \xi$, so $a/\omega^2 \sim \xi$ and $a \sim \omega^2\xi$. With this replacement, the acceleration relation becomes

$$\underbrace{\omega^2\xi}_{a} \sim g\xi k, \tag{9.68}$$

or

$$\omega^2 = 1 \times gk, \tag{9.69}$$

which is the longed-for dispersion relation with the correct dimensionless constant in red.

An exact calculation confirms the usual hope that the missing dimensionless constants are close to unity, or are unity. This fortune suggests that the procedures for choosing how to measure the lengths were reasonable. The derivation depended on two choices:

1. Replacing an exponentially falling variation in velocity potential by a step function with size equal to the length scale of the exponential decay.
2. Taking the length of the slab to be $1/k$ instead of π/k. This choice uses only 1 radian of the cycle as the characteristic length, instead of using a half cycle or π radians. Since 1 is a more natural dimensionless number than π is, choosing 1 radian rather than π or 2π radians often improves approximations.

Both approximations are usually accurate in order-of-magnitude calculations. Rarely, however, you will get caught by a factor of $(2\pi)^6$, and wish that you had used a full cycle instead of only 1 radian.

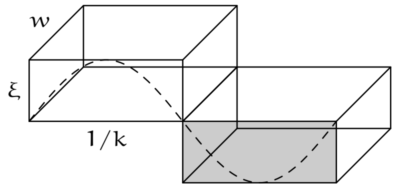

The derivation that resulted in the dispersion relation analyzed the motion of the slab using forces. Another derivation of it uses energy by balancing kinetic and potential energy. To make a wavy surface requires energy, as shown in the figure. The crest rises a characteristic height ξ above the zero of potential, which is the level surface. The volume of water moved upward is $\xi w/k$. So the potential energy is

$$\text{PE}_{\text{gravity}} \sim \underbrace{\rho\xi w/k}_{m} \times g\xi \sim \rho g w \xi^2/k. \tag{9.70}$$

The kinetic energy is contained in the sideways motion of the slab and in the upward motion of the water pushed by the slab. The slab and surface move at the same speed; they also have the same acceleration. So the sideways and upward motions contribute similar energies. If you ignore constants such as 2, you do not need to compute the energy contributed by both motions and can do the simpler computation, which is the sideways motion. The surface moves a distance ξ in a time $1/\omega$, so its velocity is $\omega\xi$. The slab has the same speed (except for constants) as the surface, so the slab's kinetic energy is

$$\text{KE}_{\text{deep}} \sim \underbrace{\rho w/k^2}_{m_{\text{slab}}} \times \underbrace{\omega^2\xi^2}_{v^2} \sim \rho\omega^2\xi^2 w/k^2. \tag{9.71}$$

This energy balances the potential energy

$$\underbrace{\rho\omega^2\xi^2 w/k^2}_{KE} \sim \underbrace{\rho g w \xi^2/k}_{PE}. \tag{9.72}$$

Canceling the factor $\rho w \xi^2$ (in red) common to both energies leaves

$$\omega^2 \sim gk. \tag{9.73}$$

The energy method agrees with the force method, as it should, because energy can be derived from force by integration. The energy derivation gives an interpretation of the dimensionless group Π_2:

$$\Pi_2 \sim \frac{\text{kinetic energy in slab}}{\text{gravitational potential energy}} \sim \frac{\omega^2}{gk}. \tag{9.74}$$

The gravity-wave dispersion relation $\omega^2 = gk$ is equivalent to $\Pi_2 \sim 1$, or to the assertion that kinetic and gravitational potential energy are comparable in wave motion. This rough equality is no surprise because waves are like springs. In spring motion, kinetic and potential energies have equal averages, a consequence of the virial theorem.

The dispersion relation was derived in three ways: by dimensional analysis, energy, and force. Using multiple methods increases our confidence not only in the result but also in the methods. 'I have said it thrice: What I tell you three times is true.'
–Lewis Carroll, *Hunting of the Snark*.

We gain confidence in the methods of dimensional analysis and in the slab model for waves. If we study nonlinear waves, for example, where the wave height is no longer infinitesimal, we can use the same techniques along with the slab model with more confidence.

With reasonable confidence in the dispersion relation, it's time study its consequences: the phase and group velocities. The crests move at the phase velocity: $v_{\text{ph}} = w/k$. For deep-water gravity waves, this velocity becomes

$$v_{\text{ph}} = \sqrt{\frac{g}{k}}, \tag{9.75}$$

or, using the dispersion relation to replace k by ω,

$$v_{\text{ph}} = \frac{g}{\omega}. \tag{9.76}$$

Let's check upstairs and downstairs. Who knows where ω belongs, but g drives the waves so it should and does live upstairs.

In an infinite, single-frequency wave train, the crests and troughs move at the phase speed. However, a finite wave train contains a mixture of frequencies, and the various frequencies move at different speeds as given by

$$v_{\text{ph}} = \frac{g}{\omega}. \tag{9.77}$$

Deep water is dispersive. Dispersion makes a finite wave train travel with the group velocity, given by $v_{\text{g}} = \partial w/\partial k$, as explained in Section 9.3.2. The group velocity is

$$v_{\text{g}} = \frac{\partial}{\partial k}\sqrt{gk} = \frac{1}{2}\sqrt{\frac{g}{k}} = \frac{1}{2}v_{\text{ph}}. \tag{9.78}$$

So the group velocity is one-half of the phase velocity, as the result for power-law dispersion relation predicts. Within a wave train, the crests move at the phase velocity, twice the group velocity, shrinking and growing to fit under the slower-moving envelope.

An everyday consequence is that ship wakes trail the ship. A ship moving with velocity v creates gravity waves with $v_{\text{ph}} = v$. The waves combine to produce wave trains that propagate forward with the group velocity, which is only $v_{\text{ph}}/2 = v/2$. From the ship's point of view, these gravity waves travel backward. In fact, they form a wedge, and the opening angle of the wedge depends on the one-half that arises from the exponent.

9.3.6 Surfing

Let's apply the dispersion relation to surfing. Following one winter storm reported in the *Los Angeles Times* – the kind of storm that brings cries of 'Surf's up!' – waves arrived at Los Angeles beaches roughly every 18 s. How fast were the storm winds that generated the waves? Wind pushes the crests as long as they move more slowly than the wind. After a long-enough push, the crests move with nearly the wind speed. Therefore the phase velocity of the waves is an accurate approximation to the wind speed.

The phase velocity is g/ω. In terms of the wave period T, this velocity is $v_{\text{ph}} = gT/2\pi$, so

$$v_{\text{wind}} \sim v_{\text{ph}} \sim \frac{\overbrace{10\,\text{m}\,\text{s}^{-2}}^{g} \times \overbrace{18\,\text{s}}^{T}}{2 \times 3} \sim 30\,\text{m}\,\text{s}^{-1}. \tag{9.79}$$

In units more familiar to Americans, this wind speed is 60 mph, which is a strong storm: about 10 on the Beaufort wind scale ('whole gale/storm'). The wavelength is given by

$$\lambda = v_{\text{ph}} T \sim 30\,\text{m}\,\text{s}^{-1} \times 18\,\text{s} \sim 500\,\text{m}. \tag{9.80}$$

On the open ocean, the crests are separated by half a kilometer. Near shore they bunch up because they feel the bottom; this bunching is a consequence of the shallow-water dispersion relation, the topic of Section 9.3.13.

In this same storm, the waves arrived at 17 s intervals the following day: a small decrease in the period. Before racing for the equations, first check that this decrease in period is reasonable. This precaution is a sanity check. If the theory is wrong about a physical effect as fundamental as a sign – whether the period should decrease or increase – then it neglects important physics. The storm winds generate waves of different wavelengths and periods, and the different wavelengths sort themselves during the trip from the far ocean to Los Angeles. Group and phase velocity are proportional to $1/\omega$, which is proportional to the period. So longer-period waves move faster, and the 18 s waves should arrive before the 17 s waves. They did! The decline in the interval allows us to calculate the distance to the storm. In their long journey, the 18 s waves raced 1 day ahead of the 17 s waves. The ratio of their group velocities is

$$\frac{\text{velocity(18 s waves)}}{\text{velocity(17 s waves)}} = \frac{18}{17} = 1 + \frac{1}{17}. \tag{9.81}$$

so the race must have lasted roughly $t \sim 17\,\text{days} \sim 1.5 \cdot 10^6\,\text{s}$. The wave train moves at the group velocity, $v_{\text{g}} = v_{\text{ph}}/2 \sim 15\,\text{m}\,\text{s}^{-1}$, so the storm distance was $d \sim t v_{\text{g}} \sim 2 \cdot 10^4\,\text{km}$, or roughly halfway around the world, an amazingly long and dissipation-free journey.

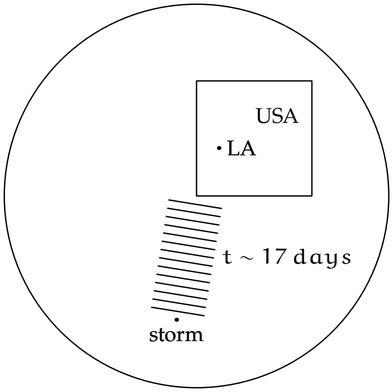

9.3.7 Speedboating

Our next application of the dispersion relation is to speedboating: How fast can a boat travel? We exclude hydroplaning boats from our analysis (even though some speedboats can hydroplane). Longer boats generally move faster than shorter boats, so it is likely that the length of the boat, l, determines the top speed. The density of water might matter. However, v (the speed), ρ, and l cannot form a dimensionless group. So look for another variable. Viscosity is irrelevant because the Reynolds number for boat travel is gigantic. Even for a small boat of length 5 m, creeping along at $2\,\mathrm{m\,s^{-1}}$,

$$\mathrm{Re} \sim \frac{500\,\mathrm{cm} \times 200\,\mathrm{cm\,s^{-1}}}{10^{-2}\,\mathrm{cm^2\,s^{-1}}} \sim 10^7. \tag{9.82}$$

At such a huge Reynolds number, the flow is turbulent and nearly independent of viscosity (Section 6.3.1). Surface tension is also irrelevant, because boats are much longer than a ripple wavelength (roughly 1 cm). The search for new variables is not meeting with success. Perhaps gravity is relevant. The four variables v, ρ, g, and l, build from three dimensions, produce one dimensionless group: v^2/gl, also called the Froude number:

$$\mathrm{Fr} \equiv \frac{v^2}{gl}. \tag{9.83}$$

The critical Froude number, which determines the maximum boat speed, is a dimensionless constant. As usual, we assume that the constant is unity. Then the maximum boating speed is:

$$v \sim \sqrt{gl}. \tag{9.84}$$

A rabbit has jumped out of our hat. What physical mechanism justifies this dimensional-analysis result? Follow the waves as a boat plows through water. The moving boat generates waves (the wake), and it rides on one of those waves. Take the bow wave: It is a gravity wave with $v_{\text{ph}} \sim v_{\text{boat}}$. Because $v_{\text{ph}}^2 = \omega^2/k^2$, the dispersion relation tells us that

$$v_{\text{boat}}^2 \sim \frac{\omega^2}{k^2} = \frac{g}{k} = g\bar{\lambda}, \tag{9.85}$$

where $\bar{\lambda} \equiv 1/k = \lambda/2\pi$. So the wavelength of the waves is roughly v_{boat}^2/g. The other length in this problem is the boat length; so the Froude number has this interpretation:

$$\text{Fr} = \frac{v_{\text{boat}}^2/g}{l} \sim \frac{\text{wavelength of bow wave}}{\text{length of boat}}. \tag{9.86}$$

Why is Fr ~ 1 the critical number, the assumption in finding the maximum boat speed? Interesting and often difficult physics occurs when a dimensionless number is near unity. In this case, the physics is as follows. The wave height changes significantly in a distance $\bar{\lambda}$; if the boat's length l is comparable to $\bar{\lambda}$, then the boat rides on its own wave and tilts upward. Tilting upward, it presents a large cross-section to the water, and the drag becomes huge. [Catamarans and hydrofoils skim the water, so this kind of drag does not limit their speed. The hydrofoil makes a much quicker trip across the English channel than the ferry makes, even though the hydrofoil is much shorter.] So the top speed is given by

$$v_{\text{boat}} \sim \sqrt{gl}. \tag{9.87}$$

For a small motorboat, with length $l \sim 5\,\text{m}$, this speed is roughly $7\,\text{m}\,\text{s}^{-1}$, or 15 mph. Boats (for example police boats) do go faster than the nominal top speed, but it takes plenty of power to fight the drag, which is why police boats have huge engines.

The Froude number in surprising places. It determines, for example, the speed at which an animal's gait changes from a walk to a trot or, for animals that do not trot, to a run. In Section 9.3.7 it determines maximum boating speed. The Froude number is a ratio of potential energy to kinetic energy, as massaging the Froude number shows:

$$\text{Fr} = \frac{v^2}{gl} = \frac{mv^2}{mgl} \sim \frac{\text{kinetic energy}}{\text{potential energy}}. \tag{9.88}$$

Here the massage technique was multiplication by unity (in red). In this example, the length l is a horizontal length, so gl is not a gravitational energy, but it has a similar structure and in other examples often has an easy interpretation as gravitational energy.

9.3.8 Walking

In the Froude number for walking speed, l is leg length, and gl is a potential energy. For a human with leg length $l \sim 1\,\text{m}$, the condition $\text{Fr} \sim 1$ implies that $v \sim 3\,\text{m}\,\text{s}^{-1}$ or 6 mph. This speed is a rough estimate for the top speed for a race walker. The world record for men's race walking was once held by Bernado Segura of Mexico. He walked 20 km in 1h:17m:25.6s, for a speed of $4.31\,\text{m}\,\text{s}^{-1}$.

This example concludes the study of gravity waves on deep water, which is one corner of the world of waves.

9.3.9 Ripples on deep water

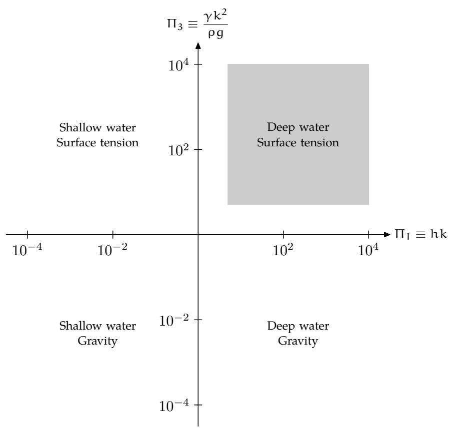

For small wavelengths (large k), surface tension rather than gravity provides the restoring force. This choice brings us to the shaded corner of the figure. If surface tension rather than gravity provides the restoring force, then g vanishes from the final dispersion relation. How to get rid of g and find the new dispersion relation? You could follow the same pattern as for gravity waves (Section 9.3.5). In that situation, the surface tension γ was irrelevant, so we discarded the group $\Pi_3 \equiv \gamma k^2/\rho g$. Here, with g irrelevant you might try the same trick: Π_3 contains g so discard it. In that argument lies infanticide, because it also throws out the physical effect that determines the restoring force, namely surface tension.

To retrieve the baby from the bathwater, you cannot throw out $\gamma k^2/\rho g$ directly. Instead you have to choose the form of the dimensionless function f_{deep} in so that only gravity vanishes from the dispersion relation.

The deep-water dispersion relation contains one power of g in front. The argument of f_{deep} also contains one power of g, in the denominator. If f_{deep} has the form $f_{\text{deep}}(x) \sim x$, then g cancels. With this choice, the dispersion relation is

$$\omega^2 = 1 \times \frac{\gamma k^3}{\rho}. \tag{9.89}$$

Again the dimensionless constant from exact calculation (in red) is unity, which we would have assumed anyway. Let's reuse the slab argument to derive this relation.

In the slab picture, replace gravitational by surface-tension energy, and again balance potential and kinetic energies. The surface of the water is like a rubber sheet. A wave disturbs the surface and stretches the sheet. This stretching creates area ΔA and therefore requires energy $\gamma \Delta A$. So to estimate the energy, estimate the extra area that a wave of amplitude ξ and wavenumber k creates. The extra area depends on the extra length in a sine wave compared to a flat line. The typical slope in the sine wave $\xi \sin kx$ is ξk. Instead of integrating to find the arc length, you can approximate the curve as a straight line with slope ξk:

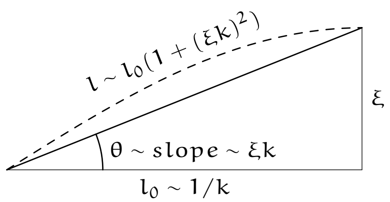

Relative to the level line, the tilted line is longer by a factor $1 + (\xi k)^2$.

As before, imagine a piece of a wave, with characteristic length $1/k$ in the x direction and width w in the y direction. The extra area is

$$\Delta A \sim \underbrace{w/k}_{\textit{level area}} \times \underbrace{(\xi k)^2}_{\textit{fractional increase}} \sim w\xi^2 k. \tag{9.90}$$

The potential energy stored in this extra surface is

$$\text{PE}_{\text{ripple}} \sim \gamma \Delta A \sim \gamma w \xi^2 k. \tag{9.91}$$

The kinetic energy in the slab is the same as it is for gravity waves, which is:

$$\text{KE} \sim \rho \omega^2 \xi^2 w / k^2. \tag{9.92}$$

Balancing the energies

$$\underbrace{\rho \omega^2 \xi^2 w / k^2}_{KE} \sim \underbrace{\gamma w \xi^2 k}_{PE}, \tag{9.93}$$

gives

$$\omega^2 \sim \gamma k^3 / \rho. \tag{9.94}$$

This dispersion relation agrees with the result from dimensional analysis. For deep-water gravity waves, we used both energy and force arguments to re-derive the dispersion relation. For ripples, we worked out the energy argument, and you are invited to work out the corresponding force argument.

The energy calculation completes the interpretations of the three dimensionless groups. Two are already done: Π_1 is the dimensionless depth and Π_2 is ratio of kinetic energy to gravitational potential energy. We constructed Π_3 as a group that compares the effects of surface tension and gravity. Using the potential energy for gravity waves and for ripples, the comparison becomes more precise:

$$\begin{aligned} \Pi_3 &\sim \frac{\text{potential energy in a ripple}}{\text{potential energy in a gravity wave}} \\ &\sim \frac{\gamma w \xi^2 k}{\rho g w \xi^2 / k} \\ &\sim \frac{\gamma k^2}{\rho g}. \end{aligned} \tag{9.95}$$

Alternatively, Π_3 compares $\gamma k^2/\rho$ with g:

$$\Pi_3 \equiv \frac{\gamma k^2 / \rho}{g}. \tag{9.96}$$

This form of Π_3 may seem like a trivial revision of $\gamma k^2/\rho g$. However, it suggests an interpretation of surface tension: that surface tension acts like an effective gravitational field with strength

$$g_{\text{surface tension}} = \gamma k^2/\rho, \tag{9.97}$$

In a balloon, the surface tension of the rubber implies a higher pressure inside than outside. Similarly in wave the water skin implies a higher pressure underneath the crest, which is curved like a balloon; and a lower pressure under the trough, which is curved opposite to a balloon. This pressure difference is just what a gravitational field with strength $g_{\text{surface tension}}$ would produce. This trick of effective gravity, which we used for the buoyant force on a falling marble (Section 6.3.2), is now promoted to a method (a trick used twice).

So replace g in the gravity-wave potential energy with this effective g to get the ripple potential energy:

$$\underbrace{\rho g w \xi^2/k}_{\textit{PE(gravity wave)}} \xrightarrow{g \to \gamma k^2/\rho} \underbrace{\gamma w \xi^2 k}_{\textit{PE(ripple)}}. \tag{9.98}$$

The left side becomes the right side after making the substitution above the arrow. The same replacement in the gravity-wave dispersion relation produces the ripple dispersion relation:

$$\omega^2 = gk \quad \xrightarrow{g \to \gamma k^2/\rho} \quad \omega^2 = \frac{\gamma k^3}{\rho}. \tag{9.99}$$

The interpretation of surface tension as effective gravity is useful when we combine our solutions for gravity waves and for ripples, in Section 9.3.11 and Section 9.3.16. Surface tension and gravity are symmetric: We could have reversed the analysis and interpreted gravity as effective surface tension. However, gravity is the more familiar force, so we use effective gravity rather than effective surface tension.

With the dispersion relation you can harvest the phase and group velocities. The phase velocity is

$$v_{\text{ph}} \equiv \frac{\omega}{k} = \sqrt{\frac{\gamma k}{\rho}}, \tag{9.100}$$

and the group velocity is

$$v_g \equiv \frac{\partial \omega}{\partial k} = \frac{3}{2} v_{ph}. \tag{9.101}$$

The factor of 3/2 is a consequence of the form of the dispersion relation: $\omega \propto k^{3/2}$; for gravity waves, $\omega \propto k^{1/2}$, and the corresponding factor is 1/2. In contrast to deep-water waves, a train of ripples moves *faster* than the phase velocity. So, ripples steam ahead of a boat, whereas gravity waves trail behind.

9.3.10 Typical ripples

Let's work out speeds for typical ripples, such as the ripples from dropping a pebble into a pond. From observation, these ripples have wavelength $\lambda \sim 1\,\mathrm{cm}$, and therefore wavenumber $k = 2\pi/\lambda \sim 6\,\mathrm{cm}^{-1}$. The surface tension of water (??) is $\gamma \sim 0.07\,\mathrm{J\,m^{-2}}$. So the phase velocity is

$$v_{ph} = \left(\frac{\overbrace{0.07\,\mathrm{J\,m^{-2}}}^{\gamma} \times \overbrace{600\,\mathrm{m^{-1}}}^{k}}{\underbrace{10^3\,\mathrm{kg\,m^{-3}}}_{\rho}} \right)^{1/2} \sim 21\,\mathrm{cm\,s^{-1}}. \tag{9.102}$$

According to relation between phase and group velocities, the group velocity is 50 percent larger than the phase velocity: $v_g \sim 30\,\mathrm{cm\,s^{-1}}$. This wavelength of 1 cm is roughly the longest wavelength that still qualifies as a ripple, as shown in an earlier figure repeated here:

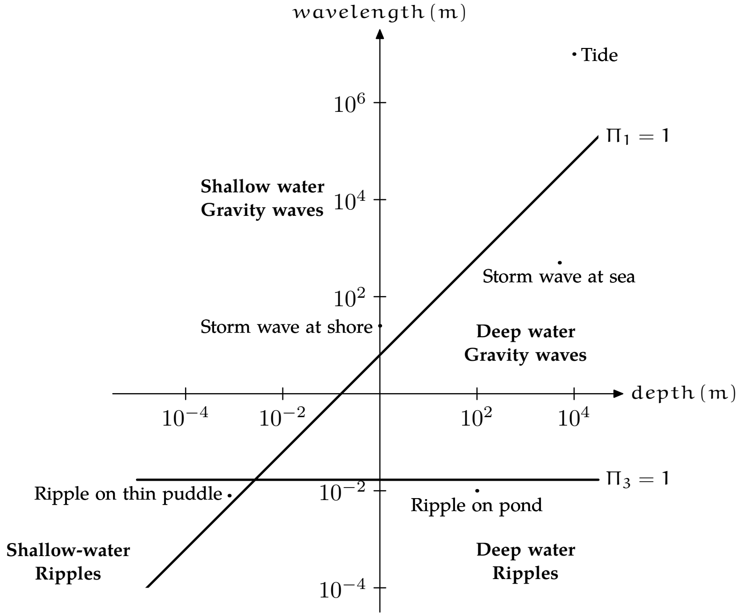

The third dimensionless group, which distinguishes ripples from gravity waves, has value

$$\Pi_3 \equiv \frac{\gamma k^2}{\rho g} \sim \frac{\overbrace{0.07\,\mathrm{J\,m^{-2}}}^{\gamma} \times \overbrace{3.6 \cdot 10^5\,\mathrm{m^{-2}}}^{k^2}}{\underbrace{10^3\,\mathrm{kg\,m^{-3}}}_{\rho} \times \underbrace{10\,\mathrm{m\,s^{-2}}}_{g}} \sim 2.6. \tag{9.103}$$

With a slightly smaller k, the value of Π_3 would slide into the gray zone $\Pi_3 \approx 1$. If k were yet smaller, the waves would be gravity waves. Other ripples, with a larger k, have a shorter wavelength, and therefore move faster: $21\,\mathrm{cm\,s^{-1}}$ is roughly the minimum phase velocity for ripples. This minimum speed explains why we see mostly $\lambda \sim 1\,\mathrm{cm}$ ripples when we drop a pebble in a pond. The pebble excites ripples of various wavelengths; the shorter ones propagate faster and the 1 cm ones straggle, so we see the stragglers clearly, without admixture of other ripples.

9.3.11 Combining ripples and gravity waves on deep water

With two corners assembled – gravity waves and ripples in deep water – you can connect the corners to form the deep-water edge. The dispersion relations, for convenience restated here, are

$$\omega^2 = \begin{cases} gk, & \text{gravity waves;} \\ \gamma k^3/\rho, & \text{ripples.} \end{cases} \quad (9.104)$$

With a little courage, you can combine the relations in these two extreme regimes to produce a dispersion relation valid for gravity waves, for ripples, and for waves in between.

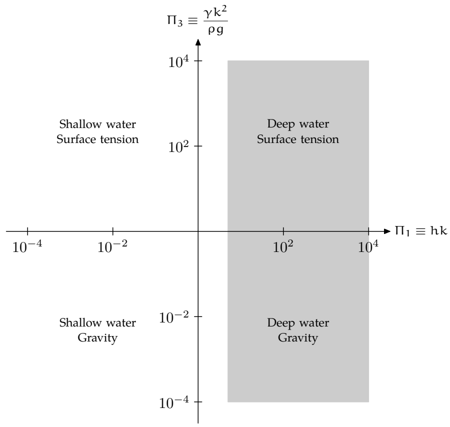

Both functional forms came from the same physical argument of balancing kinetic and potential energies. The difference was the source of the potential energy: gravity or surface tension. On the top half of the world of waves, surface tension dominates gravity; on the bottom half, gravity dominates surface tension. Perhaps in the intermediate region, the two contributions to the potential energy simply add. If so, the combination dispersion relation is the sum of the two extremes:

$$\omega^2 = gk + \gamma k^3/\rho. \quad (9.105)$$

This result is exact (which is why we used an equality). When in doubt, try the simplest solution.

You can increase your confidence in this result by using the effective gravity produced by surface tension. The two sources of gravity – real and effective – simply add, to make

$$g_{\text{total}} = g + g_{\text{surface tension}} = g + \frac{\gamma k^2}{\rho}. \quad (9.106)$$

Replace g by g_{total} in $\omega^2 = gk$ reproduces the deep-water dispersion relation:

$$\omega^2 = \left(g + \frac{\gamma k^2}{\rho}\right)k = gk + \gamma k^3/\rho. \quad (9.107)$$

This dispersion relation tells us wave speeds for all wavelengths or wavenumbers. The phase velocity is

$$v_{\mathrm{ph}} \equiv \frac{\omega}{k} = \sqrt{\frac{\gamma k}{\rho} + \frac{g}{k}}. \tag{9.108}$$

Let's check upstairs and downstairs. Surface tension and gravity drive the waves, so γ and g should be upstairs. Inertia slows the waves, so ρ should be downstairs. The phase velocity passes these tests.

As a function of wavenumber, the two terms in the square root compete to increase the speed. The surface-tension term wins at high wavenumber; the gravity term wins at low wavenumber. So there is an intermediate, minimum-speed wavenumber, k_0, which we can estimate by balancing the surface tension and gravity contributions:

$$\frac{\gamma k_0}{\rho} \sim \frac{g}{k_0}. \tag{9.109}$$

This computation is an example of order-of-magnitude minimization. The minimum-speed wavenumber is

$$k_0 \sim \sqrt{\frac{\rho g}{\gamma}}. \tag{9.110}$$

Interestingly, $1/k_0$ is the maximum size of raindrops. At this wavenumber $\Pi_3 = 1$: These waves lie just on the border between ripples and gravity waves. Their phase speed is

$$v_0 \sim \sqrt{\frac{2g}{k_0}} \sim \left(\frac{4\gamma g}{\rho}\right)^{1/4}. \tag{9.111}$$

In water, the critical wavenumber is $k_0 \sim 4\,\mathrm{cm}^{-1}$, so the critical wavelength is $\lambda_0 \sim 1.5\,\mathrm{cm}$; the speed is

$$v_0 \sim 23\,\mathrm{cm\,s}^{-1}. \tag{9.112}$$

We derived the speed dishonestly. Instead of using the maximum–minimum methods of calculus, we balanced the two contributions. A calculus derivation confirms the minimum phase velocity. A tedious calculus calculation shows that the minimum group velocity is

$$v_{\mathrm{g}} \approx 17.7\,\mathrm{cm\,s}^{-1}. \tag{9.113}$$

[If you try to reproduce this calculation, be careful because the minimum group velocity is not the group velocity at k_0.]

Let's do the minimizations honestly. The calculation is not too messy if it's done with good formula hygiene plus a useful diagram, and the proper method is useful in many physical maximum–minimum problems. We illustrate the methods by finding the minimum of the phase velocity. That equation contains constants – ρ, γ, and g – which carry through all the differentiations. To simplify the manipulations, choose a convenient set of units in which

$$\rho = \gamma = g = 1. \tag{9.114}$$

The analysis of waves uses three basic dimensions: mass, length, and time. Choosing three constants equal to unity uses up all the freedom. It is equivalent to choosing a canonical mass, length, and time, and thereby making all quantities dimensionless. Don't worry: The constants will return at the end of the minimization.

In addition to constants, the phase velocity also contains a square root. As a first step in formula hygiene, minimize instead v_{ph}^2. In the convenient unit system, it is

$$v_{\mathrm{ph}}^2 = k + \frac{1}{k}. \tag{9.115}$$

This minimization does not need calculus, even to do it exactly. The two terms are both positive, so you can use the arithmetic-mean–geometric-mean inequality (affectionately known as AM–GM) for k and $1/k$. The inequality states that, for positive a and b,

$$\underbrace{(a+b)/2}_{AM} \geq \underbrace{\sqrt{ab}}_{GM}, \tag{9.116}$$

with equality when $a = b$.

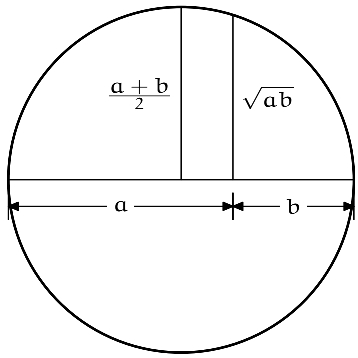

The figure shows a geometric proof of this inequality. You are invited to convince yourself that the figure is a proof. With $a = k$ and $b = 1/k$ the geometric mean is unity, so the arithmetic mean is ≥ 1. Therefore

$$k + \frac{1}{k} \geq 2, \tag{9.117}$$

with equality when $k = 1/k$, namely when $k = 1$. At this wavenumber the phase velocity is $\sqrt{2}$. Still in this unit system, the dispersion relation is

$$\omega = \sqrt{k^3 + k}, \tag{9.118}$$

and the group velocity is

$$v_{\mathrm{g}} = \frac{\partial}{\partial k}\sqrt{k^3 + k}, \tag{9.119}$$

which is

$$v_{\mathrm{g}} = \frac{1}{2}\frac{3k^2 + 1}{\sqrt{k^3 + k}}. \tag{9.120}$$

At $k = 1$ the group velocity is also $\sqrt{2}$: These borderline waves have equal phase and group velocity. This equality is reasonable. In the gravity-wave regime, the phase velocity is greater than the group velocity. In the ripple regime, the phase velocity is less than the group velocity. So they must be equal somewhere in the intermediate regime.

To convert $k = 1$ back to normal units, multiply it by unity in the form of a convenient product of ρ, γ, and g (which are each equal to 1 for the moment). How do you make a length from ρ, γ, and g? The form of the result says that $\sqrt{\rho g/\gamma}$ has units of L^{-1}. So $k = 1$ really means $k = 1 \times \sqrt{\rho g/\gamma}$, which is the same as the order-of-magnitude minimization. This exact calculation shows that the missing dimensionless constant is 1.

The minimum group velocity is more complicated than the minimum phase velocity because it requires yet another derivative. Again, remove the square root and minimize v_{g}^2. The derivative is

$$\frac{\partial}{\partial k}\underbrace{\frac{9k^4+6k^2+1}{k^3+k}}_{v_g^2} = \frac{(3k^2+1)(3k^4+6k^2-1)}{(k^3+k)^2}. \tag{9.121}$$

Equating this derivative to zero gives $3k^4+6k^2-1=0$, which is a quadratic in k^2, and has positive solution

$$k_1 = \sqrt{-1+\sqrt{4/3}} \sim 0.393. \tag{9.122}$$

At this k, the group velocity is

$$v_g(k_1) \approx 1.086. \tag{9.123}$$

In more usual units, this minimum velocity is

$$v_g \approx 1.086\left(\frac{\gamma g}{\rho}\right)^{1/4}. \tag{9.124}$$

With the density and surface tension of water, the minimum group velocity is $17.7\,\mathrm{cm\,s^{-1}}$, as claimed previously.

After dropping a pebble in a pond, you see a still circle surrounding the drop point. Then the circle expands at the minimum group velocity given. Without a handy pond, try the experiment in your kitchen sink: Fill it with water and drop in a coin or a marble. The existence of a minimum phase velocity, is useful for bugs that walk on water. If they move slower than $23\,\mathrm{cm\,s^{-1}}$, they generate no waves, which reduces the energy cost of walking.

9.3.12 Shallow water

In shallow water, the height h, absent in the deep-water calculations, returns to complicate the set of relevant variables. We are now in the shaded region of the figure. This extra length scale gives too much freedom. Dimensional analysis alone cannot deduce the shallow-water form of the magic function f in the dispersion relation. The slab argument can do the job, but it needs a few modifications for the new physical situation.

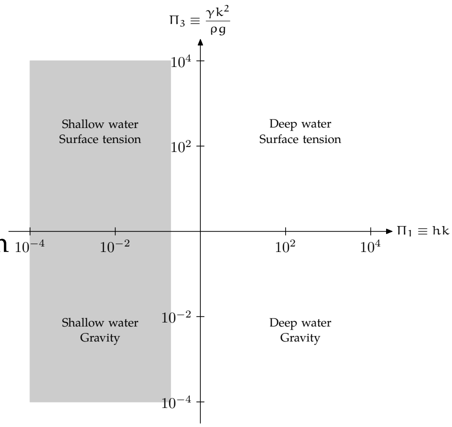

In deep water the slab has depth $1/k$. In shallow water, however, where $h \ll 1/k$, the bottom of the ocean arrives before that depth. So the shallow-water slab has depth h. Its length is still $1/k$, and its width is still w. Because the depth changed, the argument about how the water flows is slightly different. In deep water, where the slab has depth equal to length, the slab and surface move the same distance. In shallow water, with a slab thinner by hk, the surface moves more slowly than the slab because less water is being moved around. It moves more slowly by the factor hk. With wave height ξ and frequency ω, the surface moves with velocity $\xi\omega$, so the slab moves (sideways) with velocity $v_{\text{slab}} \sim \xi\omega/hk$. The kinetic energy in the water is contained mostly in the slab, because the upward motion is much slower than the slab motion. This energy is

$$\text{KE}_{\text{shallow}} \sim \underbrace{\rho wh/k}_{mass} \times \underbrace{(\xi\omega/hk)^2}_{v^2} \sim \frac{\rho w \xi^2 \omega^2}{hk^3}. \tag{9.125}$$

This energy balances the potential energy, a computation we do for the two limiting cases: ripples and gravity waves.

9.3.13 Gravity waves on shallow water

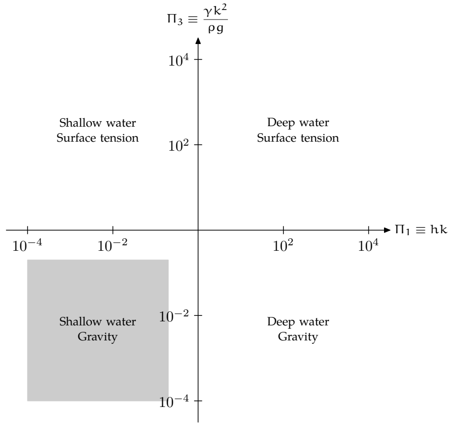

We first specialize to gravity waves – the shaded region in the figure – where water is shallow and wavelengths are long. These conditions include tidal waves,waves generated by undersea earthquakes, and waves approaching a beach. For gravity waves, the potential energy is

$$\text{PE} \sim \rho g w \xi^2 / k. \tag{9.126}$$

This energy came from the distortion of the surface, and it is the same in shallow water (as long as the wave amplitude is small compared with the depth and wavelength). [The dominant force (gravity or surface tension) determines the potential energy. As we see when we study shallow-water ripples, in Section 9.3.15, the water depth determines the kinetic energy.]

Balancing this energy against the kinetic energy gives:

$$\underbrace{\frac{\rho w \xi^2 \omega^2}{hk^3}}_{KE} \sim \underbrace{\rho g w \xi^2 / k}_{PE}. \tag{9.127}$$

So

$$\omega^2 = 1 \times ghk^2. \tag{9.128}$$

Once again, the correct, honestly calculated dimensionless constant (in red) is unity. So, for gravity waves on shallow water, the function f has the form

$$f_{\text{shallow}}(kh, \frac{\gamma k^2}{\rho g}) = kh. \tag{9.129}$$

Since $\omega \propto k^1$, the group and phase velocities are equal and independent of frequency:

$$\begin{aligned} v_{\text{ph}} &= \frac{\omega}{k} = \sqrt{gh}, \\ v_{\text{g}} &= \frac{\partial \omega}{\partial k} = \sqrt{gh}. \end{aligned} \tag{9.130}$$

Shallow water is nondispersive: All frequencies move at the same velocity, so pulses composed of various frequencies propagate without smearing.

9.3.14 Tidal waves

Undersea earthquakes illustrate the danger in such unity. If an earthquake strikes off the coast of Chile, dropping the seafloor, it generates a shallow-water wave. This wave travels without distortion to Japan. The wave speed is $v \sim \sqrt{4000\,\mathrm{m} \times 10\,\mathrm{m\,s^{-2}}} \sim 200\,\mathrm{m\,s^{-1}}$: The wave can cross a 10^4 km ocean in half a day. As it approaches shore, where the depth decreases, the wave slows, grows in amplitude, and becomes a large, destructive wave hitting land.

9.3.15 Ripples on shallow water

Ripples on shallow water – the shaded region in the figure – are rare. They occur when raindrops land in a shallow rain puddle, one whose depth is less than 1 mm. Even then, only the longest-wavelength ripples, where $\lambda \sim 1$ cm, can feel the bottom of the puddle (the requirement for the wave to be a shallow-water wave). The potential energy of the surface is given by

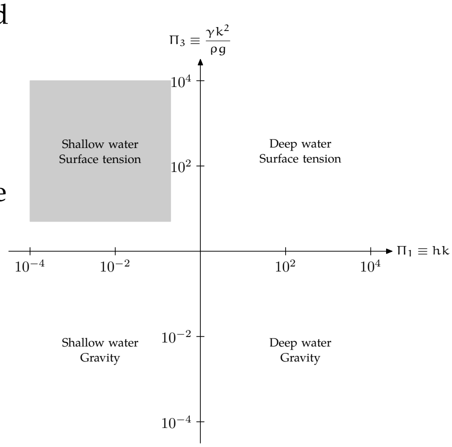

$$\mathrm{PE}_{\text{ripple}} \sim \gamma \Delta A \sim \gamma w \xi^2 k. \quad (9.131)$$

Although that formula applied to deep water, the water depth does not affect the potential energy, so we can use the same formula for shallow water.

The dominant force – here, surface tension – determines the potential energy. Balancing the potential energy and the kinetic energy gives:

$$\underbrace{\frac{\rho w \xi^2 \omega^2}{hk^3}}_{KE} \sim \underbrace{\frac{w}{k}\gamma(k\xi)^2}_{PE}. \qquad (9.132)$$

Then

$$\omega^2 \sim \frac{\gamma h k^4}{\rho}. \tag{9.133}$$

The phase velocity is

$$v_{\rm ph} = \frac{\omega}{k} = \sqrt{\frac{\gamma h k^2}{\rho}}, \tag{9.134}$$

and the group velocity is $v_{\rm g} = 2v_{\rm ph}$ (the form of the dispersion relation is $\omega \propto k^2$). For $h \sim 1\,\text{mm}$, this speed is

$$v \sim \left(\frac{0.07\,\text{N}\,\text{m}^{-1} \times 10^{-3}\,\text{m} \times 3.6 \cdot 10^5\,\text{m}^{-2}}{10^{\text{kg}}\,\text{m}^{-3}} \right)^{1/2} \sim 16\,\text{cm}\,\text{s}^{-1}. \tag{9.135}$$

9.3.16 Combining ripples and gravity waves on shallow water

This result finishes the last two corners of the world of waves: shallow-water ripples and gravity waves. Connect the corners to make an edge by studying general shallow-water waves. This region of the world of waves is shaded in the figure. You can combine the dispersion relations for ripples with that for gravity waves using two equivalent methods. Either add the two extreme-case dispersion relations or use the effective gravitational field in the gravity-wave dispersion relation. Either method produces

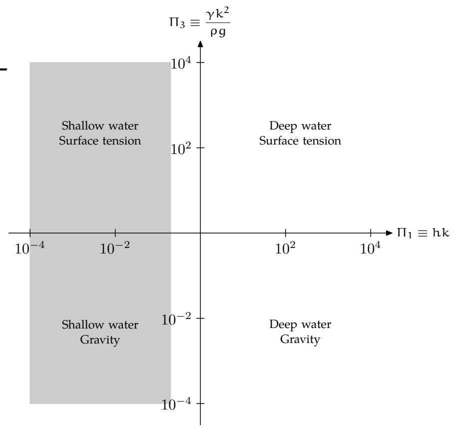

$$\omega^2 \sim k^2 \left(gh + \frac{\gamma h k^2}{\rho} \right). \tag{9.136}$$

9.3.17 Combining deep- and shallow-water gravity waves

Now examine the gravity-wave edge of the world, shaded in the figure. The deep- and shallow-water dispersion relations are:

$$\omega^2 = gk \times \begin{cases} 1, & \text{deep water;} \\ hk, & \text{shallow water.} \end{cases} \quad (9.137)$$

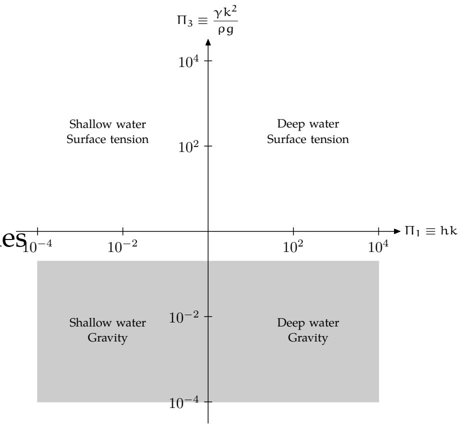

To interpolate between the two regimes requires a function $f(hk)$ that asymptotes to 1 as $hk \to \infty$ and to hk as $hk \to 0$. Arguments based on guessing functional forms have an honored history in physics. Planck derived the blackbody spectrum by interpolating between the high- and low-frequency limits of what was known at the time. We are not deriving quantum mechanics, but the principle is the same: In new areas, whether new to you or new to everyone, you need a bit of courage. One simple interpolating function is $\tanh hk$. Then the one true gravity wave dispersion relation is:

$$\omega^2 = gk \tanh hk. \quad (9.138)$$

This educated guess is plausible because $\tanh hk$ falls off exponentially as $h \to \infty$, in agreement with the argument based on Laplace's equation. In fact, this guess is correct.

9.3.18 Deep- and shallow-water ripples

We now examine the final edge: ripples in shallow and deep water, as shown in the figure. In Section 9.3.17, $\tanh kh$ interpolated between hk and 1 as hk went from 0 to ∞ (as the water went from shallow to deep). Probably the same trick works for ripples, because the Laplace-equation argument, which justified the $\tanh kh$, does not depend on the restoring force. The relevant dispersion relations:

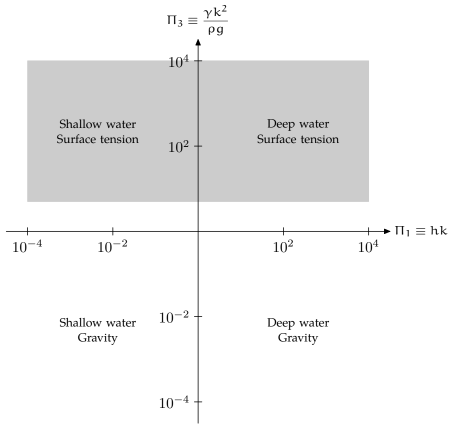

$$\omega^2 = \begin{cases} \gamma k^3/\rho, & \text{if } kh \gg 1; \\ \gamma h k^4/\rho, & \text{if } kh \ll 1. \end{cases} \tag{9.139}$$

If we factor out $\gamma k^3/\rho$, the necessary transformation becomes clear:

$$\omega^2 = \frac{\gamma k^3}{\rho} \times \begin{cases} 1, & \text{if } kh \gg 1; \\ hk, & \text{if } kh \ll 1. \end{cases} \tag{9.140}$$

This ripple result looks similar to the gravity-wave result, so make the same replacement:

$$\begin{cases} 1, & \text{if } kh \gg 1, \\ hk, & \text{if } kh \ll 1, \end{cases} \quad \text{becomes} \quad \tanh kh. \tag{9.141}$$

Then you get the general ripple dispersion relation:

$$\omega^2 = \frac{\gamma k^3}{\rho} \tanh kh. \tag{9.142}$$

This dispersion relation does not have much practical interest because, at the cost of greater complexity than the deep-water ripple dispersion relation, it adds coverage of only a rare case: ripples on ponds. We include it for completeness, to visit all four edges of the world, in preparation for the grand combination coming up next.

9.3.19 Combining all the analyses

Now we can replace g with g_{total}, to find the One True Dispersion Relation:

$$\omega^2 = (gk + \gamma k^3/\rho) \tanh kh. \tag{9.143}$$

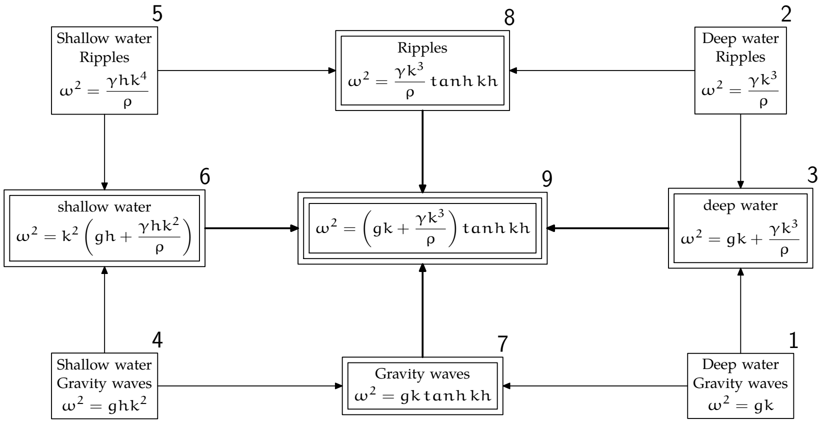

Each box in the figure represents a special case. The numbers next to the boxes mark the order in which we studied that limit. In the final step (9), we combined all the analyses into the superbox in the center, which contains the dispersion relation for all waves: gravity waves or ripples, shallow water or deep water. The arrows show how we combined smaller, more specialized corner boxes into the more general edge boxes (double ruled), and the edge regions into the universal center box (triple ruled).

In summary, we studied water waves by investigating dispersion relations. We mapped the world of waves, explored the corners and then the edges, and assembled the pieces to form an understanding of the complex, complete solution. The whole puzzle, solved, is shown in the figure. Considering limiting cases and stitching them together makes the analysis tractable and comprehensible.

9.3.20 What you have learned

1. *Phase and group velocities.* Phase velocity says how fast crests in a single wave move. In a packet of waves (several waves added together), group velocity is the phase velocity of the envelope.
2. *Discretize.* A complicated functional relationship, such as a dispersion relation, is easier to understand in a discrete limit: for example, one that allows only two (ω, k) combinations. This lumping helped explain the meaning of group velocity.

3. *Four regimes.* The four regimes of wave behavior are characterized by two dimensionless groups: a dimensionless depth and a dimensionless ratio of surface tension to gravitational energy.
4. *Look for springs.* Look for springs when a problem has kinetic- and potential-energy reservoirs and energy oscillates between them. A key characteristic of spring motion is overshoot: The system must zoom past the equilibrium configuration of zero potential energy.
5. *Most missing constants are unity.* In analyses of waves and springs, the missing dimensionless constants are usually unity. This fortunate result comes from the virial theorem, which says that the average potential and kinetic energies are equal for a $F \propto r$ force (a spring force). So balancing the two energies is exact in this case.
6. *Minimum speed.* Objects moving below a certain speed (in deep water) generate no waves. This minimum speed is the result of cooperation between gravity and surface tension. Gravity keeps long-wavelength waves moving quickly. Surface tension keeps short-wavelength waves moving quickly.
7. *Shallow-water gravity waves are non-dispersive.* Gravity waves on shallow water (which includes tidal waves on oceans!) travel at speed $\sqrt{gh}$, independent of wavelength.
8. *Froude number.* The Froude number, a ratio of kinetic to potential energy, determines the maximum speed of speedboats and of walking.

9.4 Precession of planetary orbits

Problem 9.2 AM–GM

Prove the arithmetic mean–geometric mean inequality by another method than the circle in the text. Use AM–GM for the following problem normally done with calculus. You start with a unit square, cut equal squares from each corner, then fold the flaps upwards to make a half-open box. How large should the squares be in order to maximize its volume?

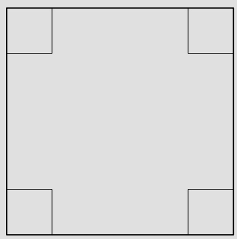

Problem 9.3 Impossible

How can tidal waves on the ocean (typical depth $\sim 4\,\text{km}$) be considered shallow water?

Problem 9.4 Oven dish

Partly fill a rectangular glass oven dish with water and play with the waves. Give the dish a slight slap and watch the wave go back and forth. How does the wave speed time vary with depth of water? Does your data agree with the theory in this chapter?

Problem 9.5 Minimum-wave speed

Take a toothpick and move it through a pan of water. By experiment, find the speed at which no waves are generated. How well does it agree with the theory in this chapter?

Problem 9.6 Kelvin wedge

Show that the opening angle in a ship wake is $2\sin^{-1}(1/3)$.

Problem 9.7 Semitones

Estimate 1.5^{14} using semitones and compare with the exact value.

Problem 9.8 Blackbody temperature of the earth

The earth's surface temperature is mostly due to solar radiation.

The solar flux $S \approx 1350\,\mathrm{W\,m^{-2}}$ is the amount of solar energy reaching the top of the earth's atmosphere. But that energy is spread over the surface of a sphere, so $S/4$ is the relevant flux for calculating the surface temperature. Some of that energy is reflected back to space by clouds or ocean before it can heat the ground, so the heating flux is slightly lower than $S/4$. A useful estimate is $S' \sim 250\,\mathrm{W\,m^{-2}}$.

Look up the Stefan–Boltzmann law (or see Problem 5.13) and use it to find the blackbody temperature of the earth.

Your value should be close to room temperature but enough colder to make you wonder about the discrepancy. Why is the actual average surface temperature warmer than the value calculated in this problem?

Problem 9.9 Xylophone notes

If you double the width, thickness, and length of a xylophone slat, what do you do to the frequency of the note that it makes?

Backmatter

Long-lasting learning

The theme of this book is how to understand new fields, whether the field is known generally but is new to you; or the field is new to everyone. In either case, certain ways of thinking promote understanding and long-term learning. This afterword illustrates these ways by using an example that has appeared twice in the book – the volume of a pyramid.

Remember nothing!

The volume is proportional to the height, because of the drilling-core argument. So $V \propto h$. But a dimensionally correct expression for the volume needs two additional lengths. They can come only from b^2. So

$$V \sim bh^2.$$

But what is the constant? It turns out to be 1/3.

Connect to other problems

Is that 3 in the denominator new information to remember? No! That piece of information also connects to other problems.

First, you can derive it by using special cases, which is the subject of Section 6.1.

Second, 3 is also the dimensionality of space. That fact is not a coincidence. Consider the simpler but analogous problem of the area of a triangle. Its area is

$$A = \frac{1}{2}bh.$$

The area has a similar form as the volume of the pyramid: A constant times a factor related to the base times the height. In two dimensions the constant is 1/2. So the 1/3 is likely to arise from the dimensionality of space.

That analysis makes the 3 easy to remember and thereby the whole formula for the volume.

But there are two follow-up questions. The first is: Why does the dimensionality of space matter? The special-cases argument explains it because you need pyramids for each direction of space (I say no more for the moment until we do the special-cases argument in lecture!).

The second follow-up question is: Does the 3 occur in other problems and for the same reason? A related place is the volume of a sphere

$$V = \frac{4}{3}\pi r^3.$$

The ancient Greeks showed that the 3 in the 4/3 is the same 3 as in the pyramid volume. To explain their picture, I'll use method to find the area of a circle then use it to find the volume of a sphere.

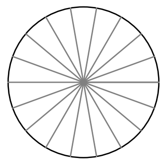

Divide a circle into many pie wedges. To find its area, cut somewhere on the circumference and unroll it into this shape:

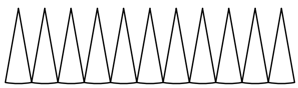

Each pie wedge is almost a triangle, so its area is $bh/2$, where the height h is approximately r. The sum of all the bases is the circumference $2\pi r$, so $A = 2\pi r \times r/2 = \pi r^2$.

Now do the same procedure with a sphere: Divide it into small pieces that are almost pyramids, then unfold it. The unfolded sphere has a base area of $4\pi r^2$, which is the surface area of the sphere. So the volume of all the mini pyramids is

$$V = \frac{1}{3} \times \underbrace{\text{height}}_{r} \times \underbrace{basearea}_{4\pi r^2} = \frac{4}{3}\pi r^3.$$

Voilà! So, if you remember the volume of a sphere – and most of us have had it etched into our minds during our schooling – then you know that the volume of a pyramid contains a factor of 3 in the denominator.

Percolation model

The moral of the preceding examples is to build connections. A physical illustration of this process is *percolation*. Imagine how oil diffuses through rock. The rock has pores through which oil moves from zone to zone. However, many pores are blocked by mineral deposits. How does the oil percolate through that kind of rock?

That question has led to an extensive mathematics research on the following idealized model. Imagine an infinite two-dimensional lattice. Now add bonds between neighbors (horizontal or vertical, not diagonal) with probability p_{bond}. The figure shows an example of a finite subsection of a percolation lattice where $p_{\text{bond}} = 0.4$. Its largest cluster – the largest set of points connected to each other – is marked in red, and contains 13% of the points.

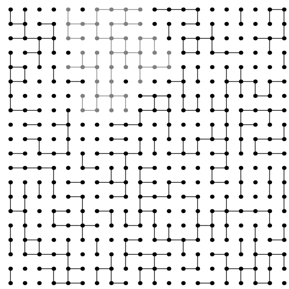

cluster = 11%
$p = 0.40$

Here is what happens as p_{bond} increases from 0.40 to 0.50 to 0.55 to 0.60:

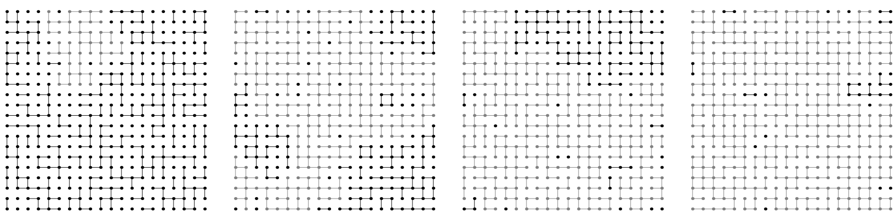

cluster = 11% $p = 0.40$ | cluster = 60% $p = 0.50$ | cluster = 70% $p = 0.55$ | cluster = 92% $p = 0.60$

The largest cluster occupies more and more of the lattice.

For an infinite lattice, a similar question is: What is the probability p_∞ of finding an infinite connected sublattice? That probability is zero until p_{bond} reaches a critical probability p_c. The critical probability depends on the topology (what kind of lattice and how many dimensions) – for the two-dimensional square lattice, $p_c = 1/2$ – but its existence is independent of topology. When $p_{\text{bond}} > p_c$, the probability of a finding an infinite lattice becomes nonzero and eventually reaches 1.0.

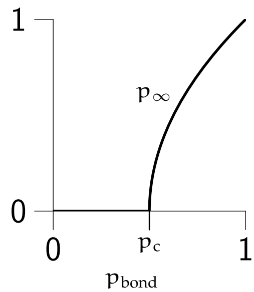

An analogy to learning is that each lattice point (each dot) is a fact or formula, and each bond links two facts. For long-lasting learning, you want the facts to support each other via their connections. Let's say that you want the facts to become part of an infinite and therefore self-supporting lattice. However, if your textbooks or way of learning means that you just add more dots – learn just more facts – then you decrease p_{bond}, so you decrease the chance of an infinite clusters. If the analogy is more exact than I think it is, you might even eliminate infinite clusters altogether.

The opposite approach is to ensure that, with each fact, you create links to facts that you already know. In the percolation model, you add bonds between the dots in order to increase p_{bond}. A famous English writer gave the same advice about life that I am giving about learning:

> Only connect! That was the whole of her sermon... Live in fragments no longer! [E. M. Forster, *Howard's End*]

The ways of reasoning presented in this book offer some ways to build those connections. Bon voyage as you learn and discover new ideas and the links between them!

Bibliography

[1] Harold Abelson and Gerald Jay Sussman with Julie Sussman. *Structure and Interpretation of Computer Programs*. MIT Press, 2nd edition, 1996.

[2] Carl Bernstein and Bob Woodward. *All the President's Men*. Simon and Schuster, New York, 1974.

[3] E. Buckingham. On physically similar systems. *Physical Review*, 4(4):345–376, 1914.

[4] Albert Einstein. *The principle of relativity*. Dover Publications, New York, 1952 [1923].

[5] Albert Einstein. *Relativity: The special and the general theory*. Three Rivers Press, reprint edition, 1995.

[6] T. E. Faber. *Fluid Dynamics for Physicists*. Cambridge University Press, Cambridge, England, 1995.

[7] Richard Feynman, Robert B. Leighton and Matthew Sands. *The Feynman Lectures on Physics*. Addison-Wesley, Reading, MA, 1963.

[8] Richard P. Feynman. *QED: The Strange Theory of Light and Matter*. Princeton, 1988.

[9] Mike Gancarz. *The UNIX Philosophy*. Digital Press, 1995.

[10] Williams James. *The Principles of Psychology*, volume 2. Henry Holt, New York, 1890.

[11] Edwin T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.

[12] David Kaiser. *Drawing Theories Apart: The Dispersion of Feynman Diagrams in Postwar Physics*. University of Chicago Press, 2005.

[13] A. Karatsuba and Yu Ofman. Multiplication of many-digital numbers by automatic computers. *Doklady Akad. Nauk SSSR*, 145:293–294, 1962. Translation in *Physics-Doklady* 7:595âĂŞ-596 (1963)..

[14] Horace Lamb. *Hydrodynamics*. Dover, New York, 6th edition, 1932.

[15] David R. Lide, editor. *CRC Handbook of Chemistry and Physics: A Ready-Reference Book of Chemical and Physical Data*. CRC Press, Boca Raton, FL, 82nd edition, 1993.

[16] Niccolo Machiavelli. *The Art of War*., 1520. http://etext.library.adelaide.edu.au/m/m

[17] Anne Marchand. Impunity for multinationals. *ATTAC*, 11 September 2002.

[18] Mars Climate Orbiter Mishap Investigation Board. Phase I report. Technical Report, NASA, 1999.

[19] Reviel Netz. *The Shaping of Deduction in Greek Mathematics: A Study in Cognitive History*. Cambridge University Press, 1999.

[20] R. C. Pankhurst. *Dimensional Analysis and Scale Factors*. Chapman and Hall, London, 1964.

[21] E. M. Purcell. Life at low Reynolds number. *American Journal of Physics*, 45:3–11, 1977.

[22] Vilayanur S. Ramachandran. The artful brain. 2003 online at www.bbc.co.uk/radio4/reith2003/lecture3.shtml.

[23] Eric S. Raymond. *The Art of UNIX Programming*. Addison-Wesley, 2003.

[24] K. A. Ross and D. E. Knuth. A programming and problem solving seminar. Technical Report, Stanford University, Stanford, California, 1989.

[25] Bertrand Russell. *A History of Western Philosophy, and Its Connection with Political and Social Circumstances from the Earliest Times to the Present Day*. Simon and Schuster, New York, 1945.

[26] Knut Schmid-Nielsen. *Scaling: Why Animal Size is So Important*. Cambridge University Press, Cambridge, England, 1984.

[27] A. SchÃűnhage and V. Strassen. Schnelle multiplikation grosser zahlen. *Computing*, 7:281–292, 1971.

[28] Herbert A. Simon and Jill H. Larkin. Why a diagram is (soemtimes) worth ten thousand words. *Cognitive Science*, 11:65–99, 1987.

[29] John Simpson and Edmund Weiner, editors. *Oxford English Dictionary*. Oxford University Press, 2nd edition, 1989. online edition (2008).

[30] Edwin F. Taylor and John Archibald Wheeler. *Spacetime Physics: Introduction to Special Relativity*. W. H. Freeman, 2nd edition, 1992.

[31] Tritton. *Physical Fluid Dynamics*. OUP, Oxford, 1988.

[32] Edward Tufte. *Beautiful Evidence*. Graphics Press, 2006.

[33] US Bureau of the Census. *Statistical Abstracts of the United States: 1992*. Government Printing Office, Washington, DC, 112th edition, 1992.

Index