Back-of-the-envelope numbers

Symbol	*What*	*Value*	*Units*
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_p^{air}	specific heat of air	1	$\mathrm{J\,g^{-1}\,K^{-1}}$
c_p	… of solids/liquids	25	$\mathrm{J\,mole^{-1}\,K^{-1}}$

Handling complexity: The art of approximation

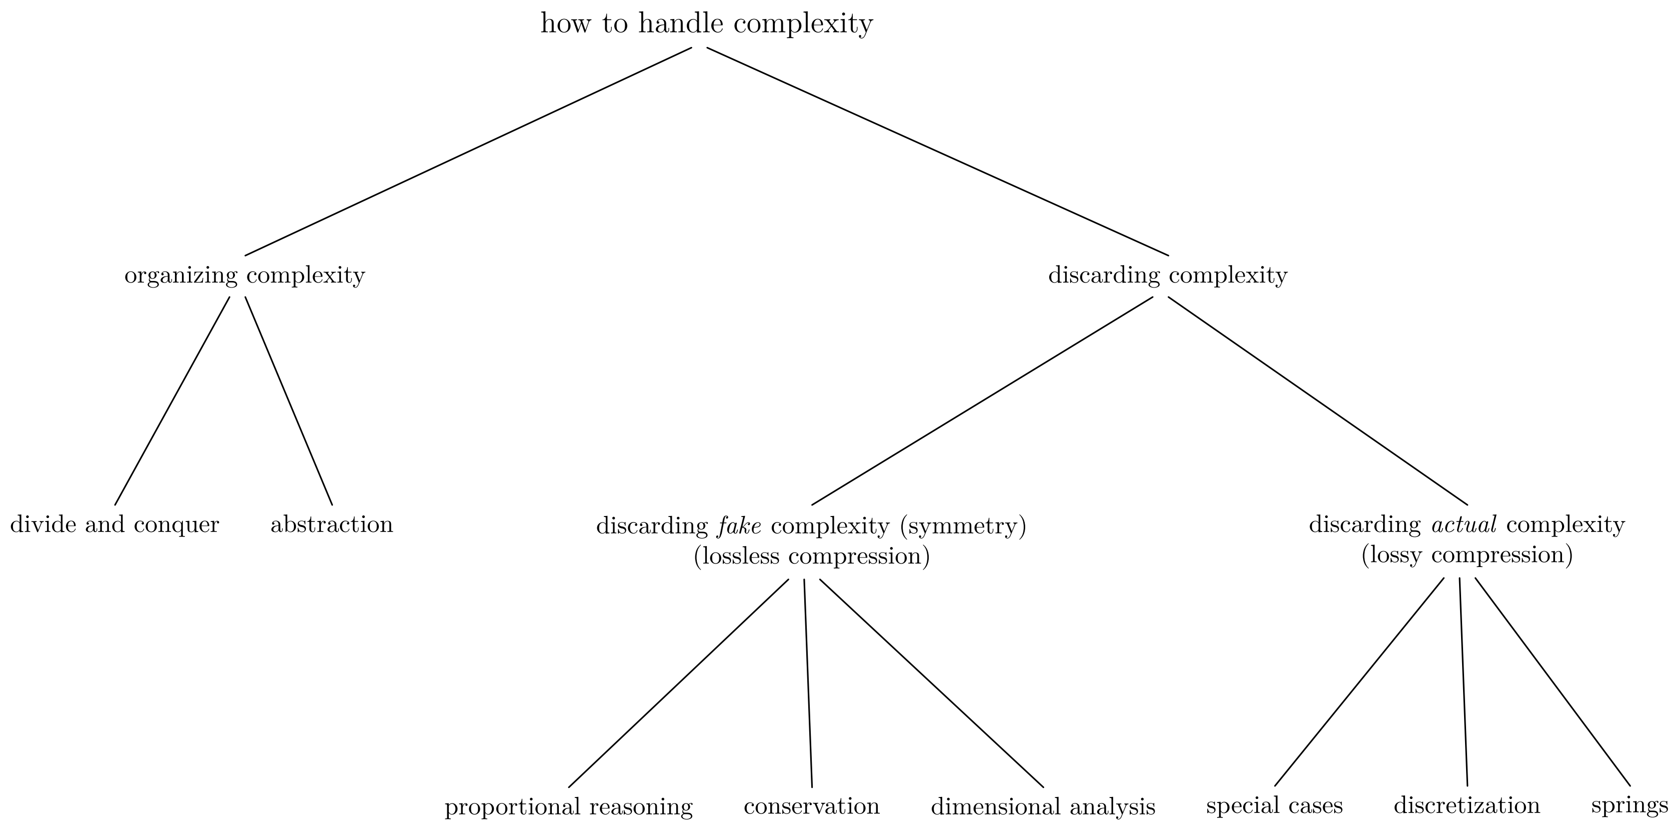

Contents

Chapter 1
Preview

An approximate model can be better than an exact model!

This counterintuitive statement suggests a few questions. First, how can approximate models be at all useful? Should we not strive for exactness? Second, what makes some models more useful than others?

On the first question: An approximate answer is all that we can understand because our minds are a small part of the world itself. So when we represent or model the world, we have to throw away aspects of the world in order for our minds to contain the model.

On the second question: Making *useful* models means discarding less important information so that our minds may grasp the important features that remains.

I wrote this book to show you how to discard the less important information and thereby to make the most useful approximate answers. From thinking about and teaching this subject for many years, I find that the most useful techniques fall into three groups:

1. **Divide and conquer** (managing complexity)
   – Hetrogenous hierarchies
   – Homogeneous hierarchies
2. **Symmetry and invariance** (removing spurious complexity)
   – Proportional reasoning
   – Conservation/box models
   – Dimensionless groups
3. **Lying** (discarding complexity)
   – Special cases
   – Spring models
   – Fractional changes
   – Discretization

The two divide-and-conquer techniques help you *manage* complexity. The three symmetry techniques help you *remove* superfluous complexity. These first two groups do lossless compression. The three lying techniques help you *discard* complexity. This third group does lossy compression.

Using these techniques, we will explore the natural and manmade worlds. Applications include:

- turbulent drag: or how falling coffee filters tell us the fuel efficiency of airplanes.
- xylophone acoustics: or why pianos are tuned with the lower notes below the ideal, equal-tempered frequency and with the higher notes above the ideal, equal-tempered frequency.
- the design of compact discs: or how Beethoven's ninth symphony helps you find the spacing between the pits.
- period of a pendulum as a function of amplitude: or how hard it was to navigate 300 years ago.
- the size distribution of eddies in turbulent flow: or how stars twinkle.
- the bending of starlight by the sun: or the size of a black hole.
- biomechanics: how high an animal jumps as a function of its size.

None of these problems has a simple analytic solution. The world – whether manmade or natural – rarely offers problems limited to one field of study, let alone problems whose equations have an analytic solution. To understand aspects of the world even partially, we need to use the preceding techniques, to make models that keep only the important features of a problem.

By making such models, we make understanding and designing more enjoyable. So the hidden although less tangible purpose of this book is to amplify your curiosity about the world.

Part 1
Divide and conquer

Divide-and-conquer reasoning – breaking large problems into small ones – is useful in many contexts. Each example of it has unique features, but two broad reasoning categories stand out. In the first category, you break the large problem into unlike, or assorted subproblems. An example is estimating the number of piano tuners in New York or, since this problem was made famous by Fermi, in Chicago, where Fermi spent much of his career. You might break it into fragments such as the number of pianos, how often each one is tuned, and how long it takes to tune a piano.

In the second category, you break the large problem into similar or identical subproblems. An example is merge sort, which breaks a list to be sorted into two halves, each sorted using merge sort – an example of recursion.

The next two chapters contain extended examples in each category.

Chapter 2

Assorted subproblems

For the first example of dividing into unlike subproblems, we estimate the spacing between pits on a CD ROM. Then we estimate the amount of oil that the United States imports annually.

2.1 Pits on a CDROM

Q: What is the spacing between the pits on a CDROM? The pits (indentations) are the memory elements, each pit storing one bit of information.

A quick estimate comes from turning over a CDROM and enjoying the brilliant colors. The colors arise because the arrangement of pits diffracts visible light by a significant angle, and the angle depends strongly on the wavelength (or color). So the pits are spaced comparably to the wavelength of light, say about 1 μm.

A second estimate might come from knowing a bit about the laser in a CD player or in a CDROM drive. It is a near-infrared laser, so its wavelength – which will be comparable to the pit size and spacing – is slightly longer than visible-light wavelengths. Since visible-light wavelengths range from 350 to 700 nm or from 0.4 to 0.8 μm, a reasonable estimate for the pit spacing is again 1 μm.

These two estimates agree, which increases our confidence in each estimate. Here is why. Because the methods are so different, an error in one method is likely to be significantly different from an error in the second method. Therefore, if the estimates agree, they are probably both reasonable. The lesson is to use as many diverse methods as you can.

The third method uses divide-and-conquer reasoning. The capacity and area together determine the pit spacing, if we make the useful approximation that the pits are regularly spaced. [This approximation is an example of discarding information, which is the extended topic of **Part 3**.]

The area is $A \sim (10\,\mathrm{cm})^2$.

The capacity is often on the box: 640 MB, which is about 5 gigabits since each byte is 8 bits. After including error-correcting bits, perhaps the capacity is 6 or 7 gigabits.

The pit spacing d comes from arranging those $N \sim 10$ gigabits into a regular lattice of bits:

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

Once again, the estimate is around $1\,\mu$m. Any result that we derive three times has to be true!

You do not need to take the capacity figure on faith. Instead, use divide-and-conquer reasoning based on how much information would be needed to encode the music on an audio CD.

The information needed depends on the play time, the sampling rate, and the sample size (number of bits).

A typical CD holds about 20 popular-music songs, each about 3 minutes long, so it is about 1 hour. An alternative piece of (perhaps bogus) history confirms this estimate: The engineers at Philips who invented the CD format and player allegedly insisted that the format hold Beethoven's Ninth Symphony, around 74 minutes.

The sampling rate is 44 kHz. Suppose you had remembered the 44 but did not remember the units: whether they were kHz or MHz. How do you choose? Human hearing extends to about 20 kHz. For comparison, the 60 Hz line-voltage hum is quite well into the audible range. Lossless sampling of sound, according to Shannon's sampling theorem, needs to have a rate of at least 2×20 kHz. The inventors of the CD format chose a slightly higher rate, so that one can make a half-decent anti-alias filter. [If you want to know more about anti-alias filters, let me know!] Even the constraint of an anti-alias filter does not require a sampling rate of 44 MHz. Indeed, the sampling rate is 44 kHz.

Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. The 16 bits is a reasonable compromise between the utopia of exact volume encoding (∞ 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 50 bits, while easy nowadays to represent digitally, implies absurd analog hardware that has an accuracy of 1 part in 2^{50}.

So the capacity is roughly

$$N \sim 1\,\text{hours} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4 \times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{2 \times 16\,\text{bits}}{1\,\text{sample}}.$$

First do the important part: the powers of ten. The 3600 contributes three; the 4.4×10^4 contributes four; and the 2×16 contributes one; for a total of eight.

The mantissas – the parts in front of the power of ten – contribute $3.6 \times 4.4 \times 3.3$. This multiplication is simplified if you remember that there are only two numbers in the world: 1 and 'few'. The only rule to remember is that ($\text{few}^2 = 10$, so 'few' acts a lot like 3. Then 3.6×3.3 is roughly 10, perhaps a bit higher. Then $3.6 \times 4.4 \times 3.3 \sim 50$.

So the estimate for the capacity is roughly $50 \times 10^8 \sim 5 \cdot 10^9$, which agrees amazingly well with the figure from a box of CDROM's. Therefore, the divide-and-conquer estimate for the capacity gives us even more confidence in our estimate for the pit spacing.

2.2 Tree representations

The structure of a divide-and-conquer estimate, with the steps of subdividing, is hierarchical. An ideal representation of hierarchical structure is a tree. Therefore to illustrate this representation, this section redoes our analysis of the pit spacing using trees.

The estimate using the area and capacity of a CDROM is the most elaborate method in **Section 2.1** for finding the pit spacing, so let's represent it as a tree. The root of the tree is 'capacity, area', a tag that reminds us of the method. As the tag suggests, to do the estimate requires finding the capacity and area, so the tree starts with two branches.

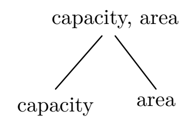

The area is easy to estimate, so the next step is to subdivide the capacity into easier parts. The first method is to look on a CDROM box, which says something like 'capacity 700MB'. A second method is to estimate the bits required to store the audio information that fit on an audio CD, by estimating the playing time, the sampling rate, and the bits per sample, where here the two channels for stereo are included in the bits per sample.

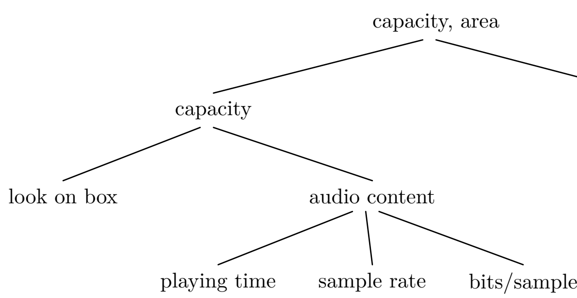

Now fill in the numbers at the leaves and propagate toward the root of the tree. The audio lasts for about an hour, which we estimated as either 20 popular music songs of 3 minutes duration or as Beethoven's Ninth Symphony. The sampling rate is 44 kHz. The samples are 32 bits each including the factor of two for stereo. The tree including these values is:

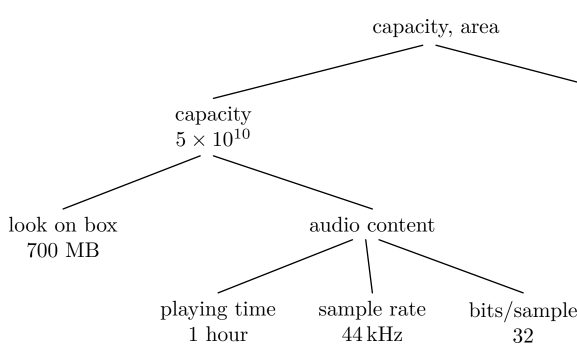

This tree is one subtree of the whole analysis. That analysis included two other methods: knowledge of the laser inside a CD player, and observing the shimmering colors due to diffraction. Including those methods – and finding that the three methods agree – makes the estimate of 1 μm robust. In pictorial terms, it makes the tree sturdy:

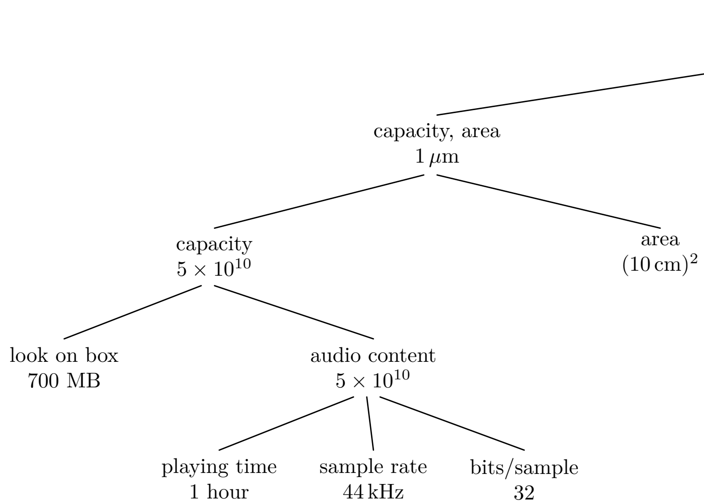

A tree is well suited for representing divide-and-conquer reasoning. This tree summarizes the whole analysis in one figure. This compact representation make it easier to find and fix mistakes in the numbers or the structure or to see which parts of the estimate are the least reliable (and probably need more subdividing).

2.3 Oil imports

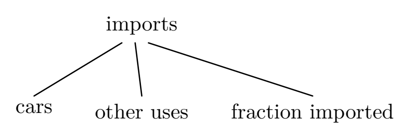

For the next example of divide-and-conquer reasoning, we will make a tree from the beginning. The problem is to estimate how much oil the United States imports, in barrels per year. There are many ways to estimate this number – good news for making robust estimates. Here I estimate it by estimating how much oil cars use, then adjusting that number to account for two items: first, that cars are not the only consumer of oil; and second, that imports are only a fraction of the oil consumed. The starting tree has just three leaves.

The rightmost two leaves are hard to guess values for, but dividing and conquering does not help. So I'll have to guess them.

Cars are a major consumer of oil, but there are other transport uses (trucks, trains, and planes), and there is heating and cooling. Given how important these other uses are, perhaps cars account for one-half of the oil consumption: a significant fraction leaving room for other significant uses. So I need to double the car result to account for other uses.

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:

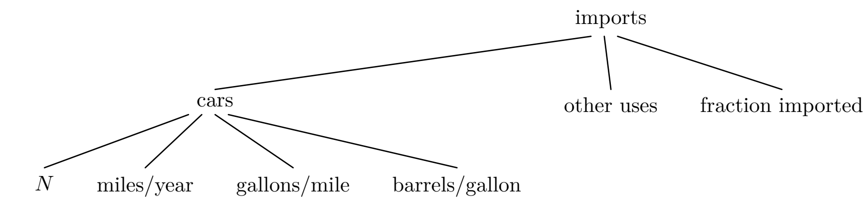

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3\times10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate 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:

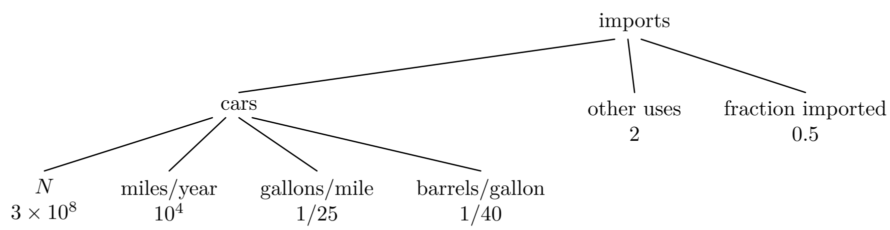

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3\times10^8\ \text{cars}\times\frac{10^4\ \text{miles}}{1\ \text{car–year}}\times\frac{1\ \text{gallon}}{25\ \text{miles}}\times\frac{1\ \text{barrel}}{40\ \text{gallons}}\sim 3\times10^9\ \text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2\times0.5=1$, so the import estimate is

$$3\times10^9\ \text{barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year!

This result, like the pit spacing, is surprisingly accurate. Why? **Section 2.5** explains a random-walk model for it, which suggests that the more you subdivide, the better.

But before discussing that model, try one more example.

2.4 Gold or bills?

Having broken into a bank vault, should you take the $100 bills or the gold?

The choice depends on how easily and losslessly you can fence the loot and on other issues outside the scope of this book. But we can study one question: Which choice lets you carry away the most money? The weight or the volume may limit how much you can carry and, more importantly for this problem, affect your choice. To make a start, let's assume that you are limited by the weight (actually, the mass) that you can carry, The problem then depends on two subproblems: the value per mass for $100 bills and for gold. In tree form:

The value per mass of gold might be a familiar figure from the newspaper or from the financial section of the evening news. It is now (2008) about \$800/oz (oz being the abbreviation for an ounce). As a rough check on the memory – e.g. should the price be \$80/oz or \$8000/oz? – here is another method. When the gold standard was reintroduced as the dollar standard in 1945, gold was set at \$35/oz. Inflation has probably devalued the dollar by at least a factor of 10 since then, so gold should be around \$350/oz now. The half-remembered figure of \$800/oz seems reasonable.

Finding the value per mass of a dollar bill starts with this tree:

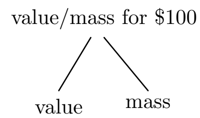

The value is specified in the problem as \$100, but the mass needs work. It breaks into the volume times the density, so the value per mass tree becomes:

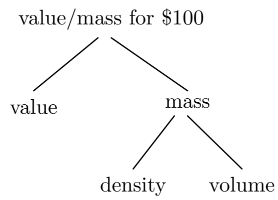

The volume breaks into length times width times thickness, so the tree grows:

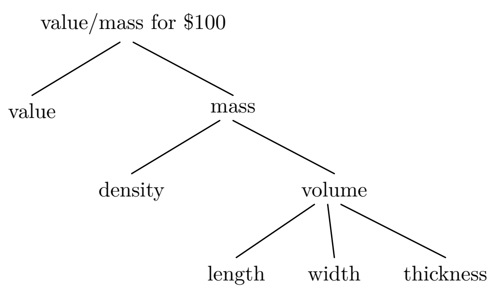

To find the length and width of a bill, lay a ruler next to a dollar bill or guess that a bill measures 2 or 3 inches by 6 inches or 6 cm × 15 cm. To develop a feel for sizes, make a guess and then, if you feel uneasy, check your answer with a ruler. As your feel for sizes develops, you will need the ruler less frequently.

Guessing the thickness of a bill is harder than guessing the length or width. However, as George Washington Plunkitt, onetime boss of Tammany Hall, said: 'I seen my opportunities and I took 'em.' Pretend that a dollar bill is made from ordinary paper. To find its thickness, look around. Next to the computer used to compose this book sits an inkjet printer; next to the printer is a ream of printer paper. If we know how thick the ream is and how many sheets it has, then we know the thickness of one sheet. You might call this technique multiply and conquer. The general lesson is that tiny values, those much below typical human experience, need to be magnified to make them easy to estimate. Large values, those much above typical human experience, need to be broken into smaller parts to make them easy to estimate. With this last step of magnifying the sheet's thickness, the full tree for the value per mass of the bill becomes:

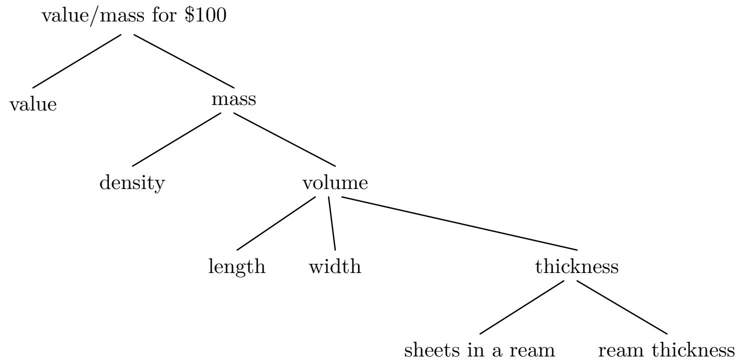

The ream (500 sheets) is roughly 5 cm thick. The only missing leaf value is the density of a bill. To find the density, use what you know: Money is paper. Paper is wood or fabric, except for many complex processing stages whose analysis is far beyond the scope of this book. When a process, here papermaking, looks formidable, forget about it and hope that you'll be okay anyway. More important is to get an estimate; correct the egregiously inaccurate assumptions later (if ever). How dense is wood? Wood barely floats, so its density is roughly that of water, which is $\rho \sim 1\,\mathrm{g\,cm^{-3}}$. So the density of a $100 bill is roughly $1\,\mathrm{g\,cm^{-3}}$.

Here is a tree including all the leaf values:

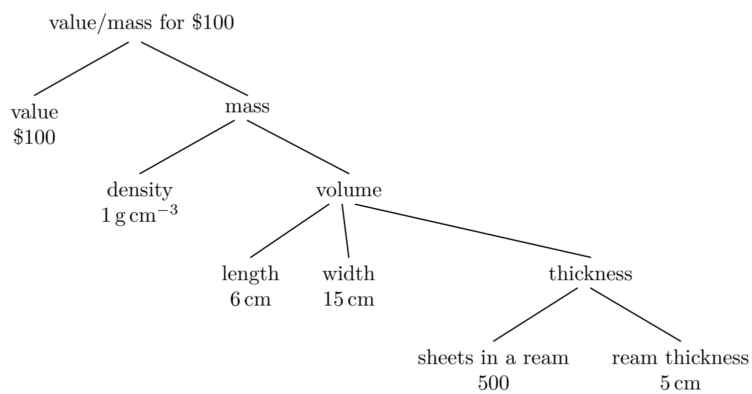

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm} \sim 1\,\mathrm{cm}^3.$$

So the mass is

$$m \sim 1\,\mathrm{cm}^3 \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}.$$

How simple! Therefore the value per mass of a $100 bill is $100/g. To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately our figure for gold is in dollars per ounce rather than per gram. Fortunately one ounce is roughly 27 g so $800/oz is roughly $30/g. Moral: Take the $100 bills but leave the $20 bills.

2.5 Random walks

The estimates in **Section 2.1** and **Section 2.3** are surprisingly accurate. The true pit spacing in a CDROM varies from 1 μm to 3 μm, according to the so-called *Red Book* where Philips and Sony give the specification of the CDROM; our estimate of 1 μm is not too bad. The true value for the oil imports is only 10% different from our estimate.

Equally important, the estimates are more accurate after doing divide-and-conquer reasoning. My 95% probability interval for oil imports, if I had to guess a value without subdividing the problem, is say from 10^6 b/yr to 10^{12} b/yr. In other words, if someone had claimed that the value is 10 million barrels per year, it would have seemed low, but I wouldn't have bet too much against it. After doing the divide-and-conquer estimate, I'd have been surprised if the true answer were more than a factor of 10 smaller or larger than the estimate.

This section presents a model for guessing in order to explain how divide-and-conquer reasoning can make estimates more accurate. The idea is that when we guess a value far outside our intuitive experience – for example, micron-sized distances or gigabarrels – the error in the exponent will be proportional to the exponent. For example, when guessing a quantity like 10^9 in one gulp, I really mean: 'It could be, say, 10^6 on the low side or, say, 10^{12} on the high side.' And when guessing a quantity like 10^{30} (the mass of the sun in kilograms), I would like to hedge my bets with a region like 10^{20} to 10^{40}. So, in this model any quantity 10^β is really shorthand for

$$10^\beta \to 10^{\beta-\beta/3} \ldots 10^{\beta+\beta/3}.$$

Now further simplify the model: Replace the range of values by its endpoints. So, if we try to guess a quantity whose true value is 10^β, we are equally likely to guess $10^{2\beta/3}$ or $10^{4\beta/3}$. A more realistic model would include 10^β as a likely possibility, but the simplest model is easy to simulate and to reason with (that justification is a fancy way to say that I am lazy).

To see the consequences of the model, I'll compare subdividing and not subdividing by using a numerical example. Suppose that we want to guess a quantity whose true value is 10^{12}. Without subdividing, we might guess 10^8 or 10^{16} (adding or subtracting one-third of the exponent), a wide range.

Compare that range to the range when we subdivide the estimate into 16 equal factors. Each factor is $10^{12/16} = 10^{3/4}$. When guessing each factor, the model says that we would guess $10^{1/2}$ or 10^1 each with $p = 0.5$. Here is an example of choosing 16 such factors randomly from $10^{1/2}$ and 10^1 and multiplying them:

$$10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5} = 10^{10.5}$$

Here are three other randomly generated examples:

$$10^{1}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5} = 10^{13.0}$$
$$10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5} = 10^{11.5}$$
$$10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{1}\cdot 10^{1}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{0.5}\cdot 10^{1}\cdot 10^{0.5} = 10^{10.5}$$

These estimates are mostly within one factor of 10 from the true answer of 10^{12}, whereas the one-shot estimate might be off by four factors of 10. What has happened is that the errors in the individual pieces are unlikely to point in the same direction. Some pieces will be underestimates, some will be overestimates, and the product of all the pieces is likely to be close to the true value.

This numerical example is our first experience with the random walk. Their crucial feature is that the expected wanderings are significantly smaller than if one walks in a straight line without switching back and forth. How much smaller is a question that we will answer in **Chapter 8** when we introduce special-cases reasoning.

2.6 The Unix philosophy

Organizing complexity by breaking it into manageable parts is not limited to numerical estimation; it is a general design principle. It pervades the Unix and its offspring operating systems such as GNU/Linux and FreeBSD. This section discusses a few examples.

2.6.1 Building blocks and pipelines

Here are a few of Unix's building-blocks programs:

- `head`: prints the first n lines from the input; for example, `head -15` prints the first 15 lines.
- `tail`: prints the last n lines from the input; for example, `tail -15` prints the last 15 lines.

How can you use these building blocks to print the 23rd line of a file? Divide and conquer! One solution is to break the problem into two parts: printing the first 23 lines and, from those lines, printing the last line. The first subproblem is solved with `head -23`. The second subproblem is solved with `tail -1`.

To combine solutions, Unix provides the pipe operator. Denoted by the vertical bar |, it connects the output of one program to the input of another command. In the numerical estimation problems, we combined the solutions to the subproblems by using multiplication. The pipe operator is analogous to multiplication. Both multiplication in numerical estimation, and pipes in programming, are examples of composition operators, which are essential to a divide-and-conquer solution.

To print the 23rd line, use this combination:

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

To tell the system where to get the input, there are alternatives:

1. Use the preceding combination as is. Then the input comes from the keyboard, and the combination will read 23 typed lines, print out the final line from those 23 lines, and then will exit.
2. Tell `head` to get its input from a file. An example file is the dictionary. On my GNU/Linux laptop it is the file `/usr/share/dict/words`, with one word per line. To print the 23rd line (i.e. the 23rd word):

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

3. Let `head` read from its idea of the keyboard, but connect the keyboard to a file. This method uses the < syntax:

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

   The < operator tells the shell (the Unix command interpreter) to connect `/usr/share/dict/words` to the input of `head`.

4. Like the preceding method, but use the `cat` program. The `cat` program copies its input file(s) to the output. So this extended pipeline has the same effect as the preceding alternative:

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

   It is slightly less efficient than letting the shell redirect the input itself, because the longer pipeline requires running one extra program (`cat`).

This example introduced the Unix philosophy: To enable divide-and-conquer reasoning, provide useful small utilities and ways to combine them. The next section applies this philosophy to a whimsical example from a scavenger hunt created by Donald Knuth: Find the next word in the dictionary after 'angry', where the dictionary is alphabetized starting with the last letter, then the second-to-last letter, etc.

2.6.2 Sorting and searching

So, how do you find the next word in the dictionary after 'angry', where the dictionary is alphabetized starting with the last letter, then the second-to-last letter, etc.?

Divide the problem into two parts:

1. Make a reverse dictionary, alphabetized starting with the last letter, then the second-to-last letter, etc.
2. Printing the line after 'angry'.

The first problem subdivides into:

1. Reverse each line of a dictionary.
2. Sort the reversed dictionary.
3. Unreverse each line.

Unix provides `sort` for the second subproblem. For the first and third problems, a search through the Unix toolbox, using `man -k`, says:

```
$ man -k reverse
build-rdeps (1)      - find packages that depend on a specific package to
bui...
col (1)              - filter reverse line feeds from input
git-rev-list (1)     - Lists commit objects in reverse chronological order
rev (1)              - reverse lines of a file or files
tac (1)              - concatenate and print files in reverse
xxd (1)              - make a hexdump or do the reverse.
```

Ah! `rev` is just the program for us. So the first subproblem is solved with this pipeline:

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

The second problem – finding the line after 'angry' – is a task for the pattern-finding program `grep`. In the simplest usage, you tell `grep` a pattern, and it prints every line from its input that matches the pattern.

The patterns are regular expressions. Their syntax can become arcane, but the most important features are simple. For example,

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

prints all lines that exactly match `angry`: The ^ character matches the beginning of the line, and the $ character matches the end of the line.

That invocation of `grep` is not useful except as a spell checker, since it tells us only that `angry` is in the dictionary. However, the `-A` option, you can tell `grep` how many lines to print after each matching line. So

```
grep -A 1 '^angry$' < /usr/share/dict/words
```

will print 'angry' and the word after it (in the regular dictionary):

```
angry
angst
```

To print just the word after 'angry', follow the `grep` command with `tail`:

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

Now combine these two solutions into solving the scavenger hunt problem:

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

This pipeline fails with the error

```
rev: stdin: Invalid or incomplete multibyte or wide character
```

The rev program is complaining that it doesn't understand some of the characters in the dictionary. rev is from the old, ASCII-only days of Unix, whereas the dictionary is modern and includes non-ASCII characters such as accented letters.

To solve this unexpected problem, clean the dictionary before passing it to rev. The cleaning program is again grep, which can allow through only those lines that are pure ASCII. This command

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

will print a dictionary made up only of unaccented, lowercase letters. In a regular expression, the * operator means 'match 0 or more occurrences of the preceding regular expression'.

The full pipeline is

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

where the backslashes at the end of the lines tell the shell to keep reading the command even though the line ended.

The tree representing this solution is

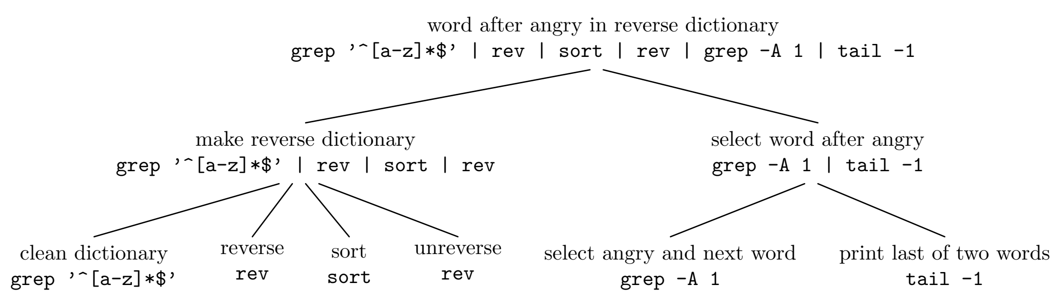

which produces 'hungry'.

2.6.3 Further reading

To learn more about the principles of Unix, especially how the design facilitates divide-and-conquer programming, see [1, 2, 3].

Chapter 3

Alike subproblems

Part 2

Symmetry and Invariance

The first part of this book discussed how to organize, and therefore how to manage complexity. The second and third parts discuss how to *eliminate complexity*, with this second part discussing three methods for finding and removing complexity that is not real.

The three methods are proportional reasoning, conservation, and dimensional analysis, and are examples of symmetry reasoning. Symmetry is also a powerful technique on its own, even without using its for the three methods. The next chapter therefore presents general examples of symmetry reasoning, and the following chapters develop this method reasoning into the three methods of proportional reasoning, conservation, and dimensional analysis.

Chapter 4
Symmetry

Symmetry is often thought of as a purely geometric concept, but it is useful in a wide variety of problems. Whenver you can use symmetry, use it and will simplify the solution. The following sections illustrate symmetry in calculus, geometry, and heat transfer.

4.1 Calculus

For what value of x is $3x - x^2$ a maximum?

The usual method is to take the derivative:

$$\frac{d}{dx}(3x - x^2) = 3 - 2x = 0,$$

whereupon $x_{\text{max}} = 3/2$.

Although differentiating is a general method, its generality comes at a cost: that its results are often hard to interpret. One does the manipulations, and whatever formulas show up at the end, so be it. So, if you can find a simplification, you are likely to get a more insight into why the answer came out the way that it did.

For this problem, symmetry simplifies it enough that nothing remains to do. To see how, first factor the equation into $x(3 - x)$. Let x_{max} be where it has its maximum. The factors x and $3 - x$ can be swapped using the substitution $x' = 3 - x$. In terms of x', the problem becomes maximizing $(3 - x')x'$. This formula has the same structure as the original one $x(3 - x)$! So the symmetry operation preserves this structure. Since the x or x' location of the maximum depends only on the structure, the location has the same numerical value whether in the x or x' coordinate systems. So it is said to be invariant under the substitution operation. Therefore, in this problem, the $x' \rightarrow 3 - x$ substitution is a symmetry.

Since $x' = 3 - x$ and, as a result of symmetry, $x'_{\text{max}} = x_{\text{max}}$, the only solution is $x_{\text{max}} = x'_{\text{max}} = 3/2$.

A similar, perhaps more telegraphic argument, is that the maximum is halfway between the two roots $x = 0$ and $x = 3$, so the maximum is, again, at $x_{\text{max}} = 3/2$. This argument implicitly contains symmetry, which is the justification for saying that the maximum is midway between the roots.

The next calculus example, from electrical and mechanical engineering, is to maximize the response of a second-order system such as a damped spring–mass system or an *LRC* circuit. The response depends on the frequency and amplitude of the driving input, and is measured as the ratio of output to input amplitude. This ratio is the gain A, and a few applications of Newton's second law produces

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

where Q is the quality factor of the system (the inverse of the damping), j is $\sqrt{-1}$ and ω is measured in units of the natural frequency.

The problem is to find the peak response, meaning the frequency ω_{max} that maximizes the magnitude of the gain and the gain at that frequency. The magnitude of the gain is

$$|A(\omega)| = \frac{\omega}{\sqrt{(1-\omega^2)^2 + \omega^2/Q^2}}$$

Because of the squares and square roots, a brute-force approach by taking the derivative will generate messy equations. So, use symmetry. What is the symmetry operation? It will be be a flip of the coordinate system, but around what point? The value $\omega = 1$ is special because that choice eliminates the denominator term $(1-\omega^2)^2$, which helps to minimize the denominator and maximize the gain. On the other hand, decreasing ω slightly could increase the gain because, at the cost of increasing $(1-\omega^2)^2$, it decreases the ω^2/Q^2 term in the denominator. On the other hand, increasing ω slightly might produce a higher gain because it increases the numerator of the gain.

To summarize: $\omega = 1$ is special but slightly higher or lower than $\omega = 1$ could be optimal too. Since $\omega = 1$ is special, use it as the point that is preserved by the symmetry operation. For a symmetry operation, interchange the $\omega < 1$ and $\omega > 1$ ranges. Frequencies mostly matter as ratios to one another – for example in music – so do the interchange by defining $\omega' = 1/\omega$ rather than $\omega' = 1 - \omega$. With the reciprocal definition, the problem becomes to maximize the magnitude of $A(\omega')$, where

$$A(\omega') = \frac{j/\omega'}{1 + j/\omega' Q - 1/\omega'^2}.$$

Multiply numerator and denominator by 1 in the form of ω'^2/ω'^2:

$$A(\omega') = \frac{j\omega'}{\omega'^2 + j\omega/Q - 1}.$$

Its magnitude is

$$|A(\omega')| = \frac{\omega'}{\sqrt{(1-\omega'^2)^2 + \omega'^2/Q^2}}.$$

This formula has the same structure as the magnitude in terms of ω itself, and this information is enough to solve for ω_{max}. Because of the isomorphic structure, $\omega'_{\text{max}} = \omega_{\text{max}}$. But by construction $\omega' = 1/\omega$, so ω'_{max} is also $1/\omega_{\text{max}}$. The only solution is $\omega_{\text{max}} = \pm 1$. Since the negative root is boring, the relevant solution is $\omega_{\text{max}} = 1$ and the response there is

$$A(\omega_{\max}) = \frac{j}{1/Q} = jQ.$$

The factor of Q in the maximum response says that a lightly damped system, where $Q \gg 1$, can reach a high amplitude if you push it at the so-called resonant frequency. The j says that the response at this resonant frequency lags the input by 90 degrees. In other words, the greatest push happens when the velocity, not the displacement, is a maximum.

4.2 Graphical symmetry

The following pictorial problem illustrates symmetry applied to a geometric problem, the traditional domain of symmetry:

> How do you cut an equilateral triangle into two equal halves using the shortest, not-necessarily-straight path?

Here are several candidates among the infinite set of possibilities for the path.

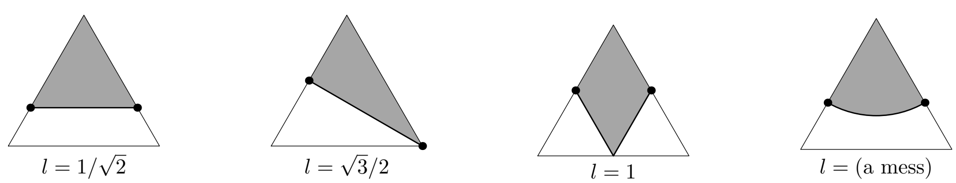

Let's compute the lengths of each bisecting path, with length measured in units of the triangle side. The first candidate encloses an equilateral triangle with one-half the area of the original triangle, so the sides of the smaller, shaded triangle are smaller by a factor of $\sqrt{2}$. Thus the path, being one of those sides, has length $1/\sqrt{2}$. In the second choice, the path is an altitude of the original triangle, which means its length is $\sqrt{3}/2$, so it is longer than the first candidate. The third candidate encloses a diamond made from two small equilateral triangles. Each small triangle has one-fourth the area of the original triangle with side length one, so each small triangle has side length $1/2$. The bisecting path is two sides of a small triangle, so its length is 1. This candidate is longer than the other two.

The fourth candidate is one-sixth of a circle. To find its length, find the radius r of the circle. One-sixth of the circle has one-half the area of the triangle, so

$$\underbrace{\pi r^2}_{A_{\text{circle}}} = 6 \times \frac{1}{2} A_{\text{triangle}} = 6 \times \frac{1}{2} \times \underbrace{\frac{1}{2} \times 1 \times \frac{\sqrt{3}}{2}}_{A_{\text{triangle}}}.$$

Multiplying the pieces gives

$$\pi r^2 = \frac{3\sqrt{3}}{4},$$

and

$$r = \sqrt{\frac{3\sqrt{3}}{4\pi}}.$$

The bisection path is one-sixth of a circle, so its length is

$$l = \frac{2\pi r}{6} = \frac{\pi}{3}\sqrt{\frac{3\sqrt{3}}{4\pi}} = \sqrt{\frac{\pi\sqrt{3}}{12}}.$$

The best previous candidate (the first picture) has length $1/\sqrt{2} = 0.707\ldots$. Does the mess of π and square roots produce a shorter path? Roll the drums. . . :

$$l = 0.67338\ldots,$$

which is less than $1/\sqrt{2}$. So the circular arc is the best bisection path *so far*. However, is it the best among all possible paths? The arc-length calculation for the circle is messy, and most other paths do not even have a closed form for their arc lengths.

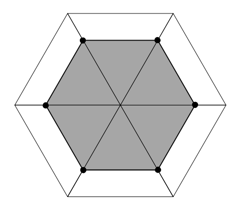

Instead of making elaborate calculations on every path, of which there are many, try **symmetry**, which is the mathematical principle for the three methods in this part of the book. To use symmetry, replicate the triangle six times to make a hexagon, thereby replicating the candidate path as well.

Here is the result of replicating the first candidate where the bisection line goes straight across. The original triangle becomes the large hexagon, and the enclosed half-triangle becomes a smaller hexagon having one-half the area of the large hexagon.

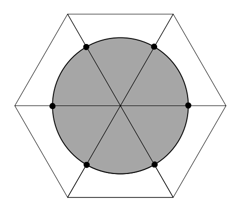

Compare that picture with the result of replicating the circular-arc bisection. The large hexagon is the same as for the last replication, but now the bisected area replicates into a circle. Which path has the shorter perimeter, the shaded hexagon or this circle? The **isoperimetric theorem** says that among all shapes with the same area the circle has the smallest perimeter. Since the circle and the smaller hexagon enclose the same area – which is three times the area of one triangle – the circle has a smaller perimeter than the hexagon, and it has a smaller perimeter than the result of replicating any other bisecting path. So the circular arc is the solution.

The lesson of this example is that symmetry can remove complexity. The complexity in this problem comes from the edges of the triangle: How much of each edge should be part of the bisected shape? Different paths use different amounts of each edge, and there's no obvious way to deduce the correct amounts. After making the figure symmetric by replicating the triangle into a hexagon, the edges become irrelevant. In the symmetric figure, the question simplifies to finding the shortest path that encloses one-half of the hexagon.

4.3 Heat flow

Here is a metal sheet in the shape of a regular pentagon with the sides held at fixed temperatures. What is the temperature at the center of the pentagon?

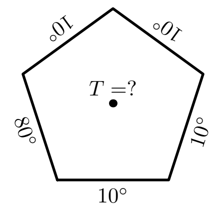

This problem is difficult to solve analytically because heat flow is described by a second-order partial differential equation, and this equation has simple solutions only for a few simple boundaries. A pentagon is, alas, not one of those boundaries. Symmetry, however, makes the solution flow.

The symmetry operation is rotation because the pentagon's orientation is an arbitrary choice. Nature, in the person of the heat equation, does not care how we point our coordinate systems. So these five orientations of the sheet behave the same:

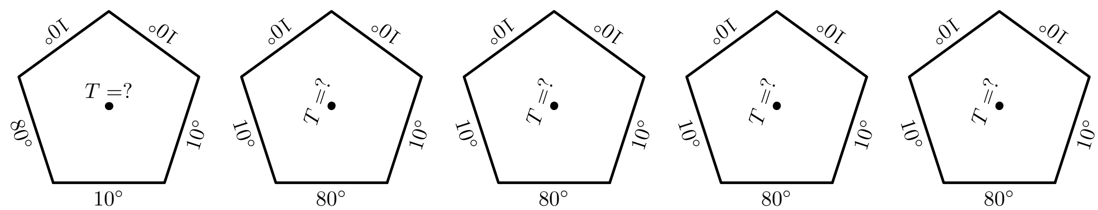

Now stack these sheets (mentally), adding their temperatures that lie on top of each other to get the temperature profile of a new metal sheet. For the new sheet, each edge has temperature

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

Therefore, the entire sheet is at 120°.

Since the symmetry operation is a rotation (by 72°) about the center, the centers overlap when the plates are (mentally) stacked one on top of the other. Furtherfore, heat flow is proportional to temperature difference – i.e. heat flow is a linear process – so the temperature in the interior of the combined plate is the sum of the five corresponding interior temperatures. Since the stacked plate has a temperature of 120° throughout it, and the centers of the five subsheets align on top of each other, each center is at $T = 120^\circ/5 = 24^\circ$.

4.4 Looking forward

The next three chapters use this aspect of symmetry – finding and removing fake complexity – to develop three techniques: proportional reasoning, conservation, and dimensional analysis.

Chapter 5
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.

5.1 Period of a spring–mass system

As a first example of proportional reasoning, here is one way to explain a famous result in physics: that the period of spring–mass system is independent of the amplitude.

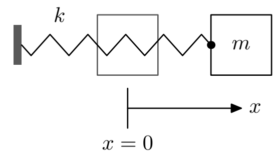

So imagine a mass m connected to the wall by a spring with spring constant k. If disturbed, the mass oscillates. The period of the system is the time for the mass to make a round trip through the equilibrium position.

Extend the spring by a distance x_0; this displacement is the amplitude. To see how it affects the period, make an approximation, which will be an example of throwing away information (the topic of **Part 3**). The approximation is to pretend that the pendulum moves with a constant speed v. Then the period is

$$T \sim \frac{\text{distance}}{\text{speed } v},$$

and the distance that the mass travels in one period is $4x_0$. Ignore the factor of 4:

$$T \sim \frac{x_0}{v}.$$

Proportional reasoning helps us estimate v by an energy argument. The initial potential energy is $\text{PE} \sim kx_0^2$ or

$$\text{PE} \propto x_0^2.$$

The maximum kinetic energy, which we use as a proxy for the typical kinetic energy, is the initial potential energy, so

$$\text{KE}_{\text{typical}} \propto x_0^2$$

as well. The typical velocity is $\sqrt{\text{KE}_{\text{typical}}}$, so

$$v_{\text{typical}} \propto x_0.$$

That result is great news because it means that the period is proportional to 1:

$$T \propto \frac{x_0}{x_0} = x_0^0.$$

In other words, the period is independent of amplitude.

5.2 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 **Part 3**. 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:

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

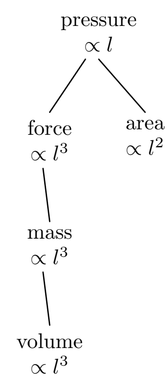

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

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.

5.3 Animal jump heights

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.

5.3.1 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.

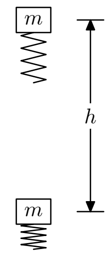

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{\text{jump}} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just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

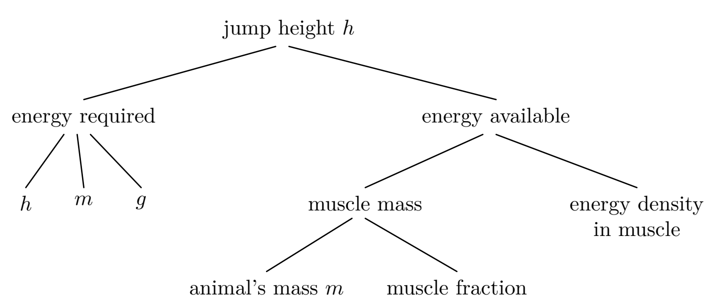

Now finish propagating toward the root. The available energy is

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

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* [4, 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

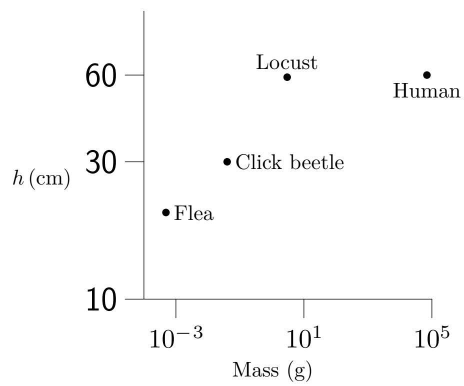

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

5.3.2 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.

5.4 Drag

This section section contains a proportional-reasoning analysis of drag – using a home experiment – and then applies the results to jumping fleas.

5.4.1 Home experiment using falling cones

Here is a home experiment for understanding drag. Photocopy this page and cut out these templates, then tape the edges together to make a cone:

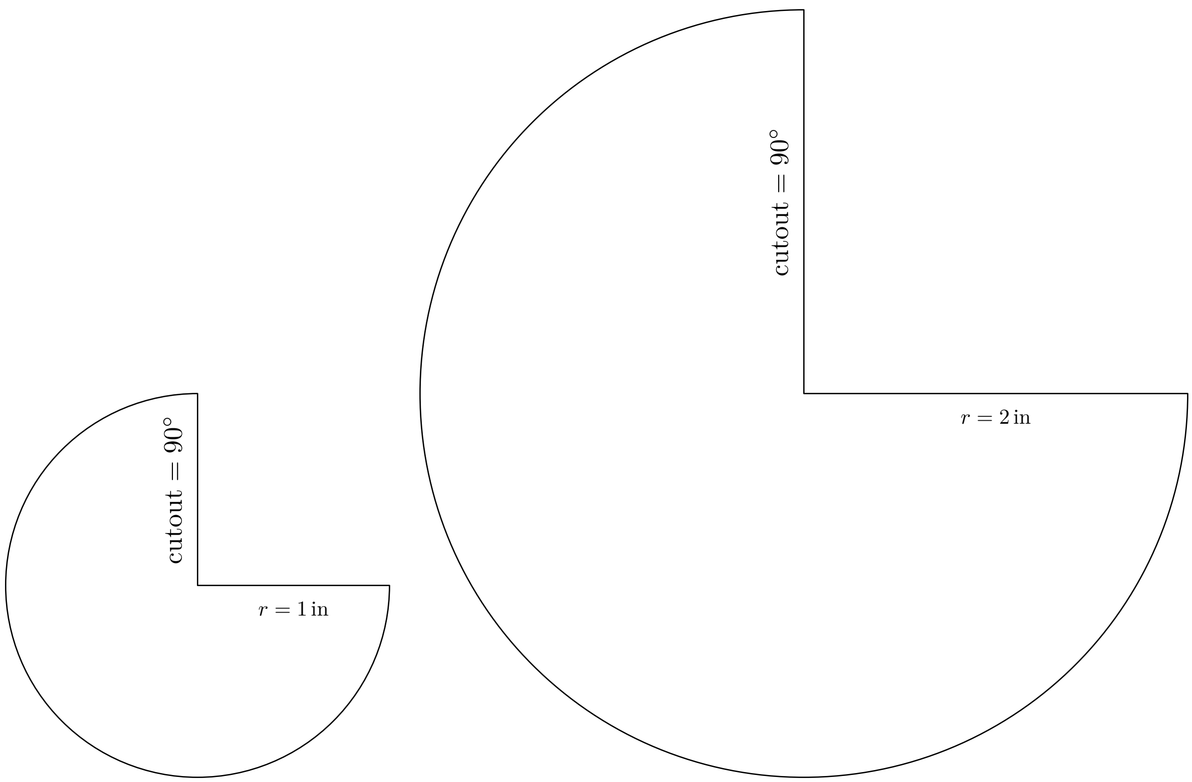

If you drop the small cone and the big cone, which falls faster? In particular, what is the ratio of their fall times $t_{\text{big}}/t_{\text{small}}$? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) 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' – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Proportional reasoning, supplemented with judicious lying, is a simple and quick alternative.

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

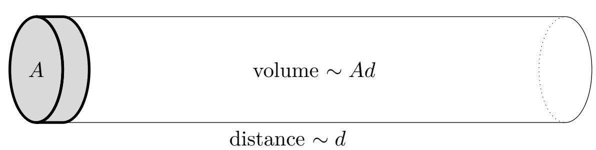

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is ρAd, where ρ is the fluid density (here, the air density). The 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:

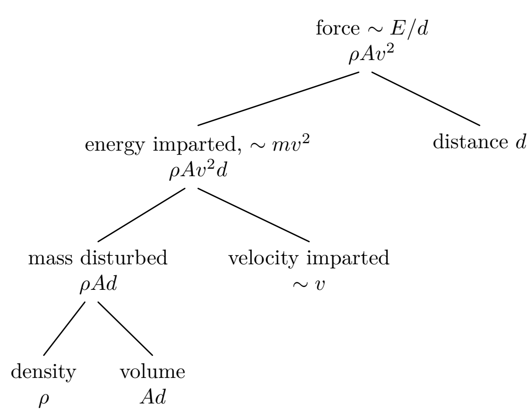

The result that $F_{\text{drag}} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly – a result discussed later in the book in **Part 3** – so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

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

Since the air density ρ is the same for the large and small cone, the relation simplifies to

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

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

$$v \propto \sqrt{\frac{W}{A}}.$$

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. With $W \propto A$, the terminal velocity becomes

$$v \propto \sqrt{\frac{A}{A}} = A^0.$$

In other words, the terminal velocity is independent of A, so the small and large cones should fall at the same speed. To test this prediction, I stood on a handy table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another!

5.4.2 Effect of drag on fleas jumping

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

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

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 $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{\text{animal}} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. 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 $l \sim 10^{-3}$ m. So

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

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.

5.4.3 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 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 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:

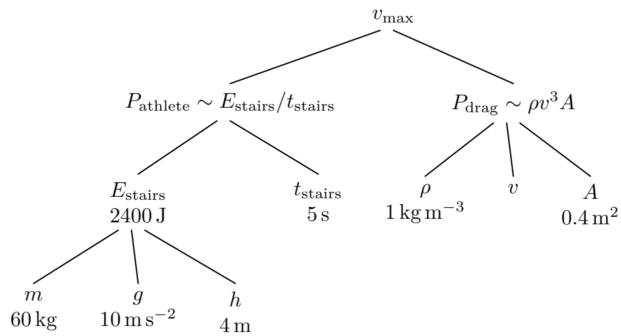

Now combine the estimates to find the maximum speed. Putting in numbers gives

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

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:

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

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

5.4.4 Swimming

The last section's analysis of cycling helps predict the world-record speed for swimming. The last section showed that

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

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 $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

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

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 $1.8\,\text{m}\,\text{s}^{-1}$.

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!]

5.4.5 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 $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ 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.

5.5 Analysis of algorithms

Proportional reasoning is the basis of an entire subject of the analysis of algorithms, a core part of computer science. How fast does an algorithm run? How much space does it require? A proportional-reasoning analysis helps you decide which algorithms to use. This section discusses these decisions using the problem of how to square very large numbers.

Squaring numbers is a special case of multiplication, but the algebra is simpler for squaring than for multiplying since having only one number as the input means there are fewer variables in the analysis.

Here is a divide-and-conquer version of the standard school multiplication algorithm.

Chapter 6

Box models and conservation

6.1 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:

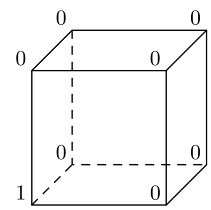

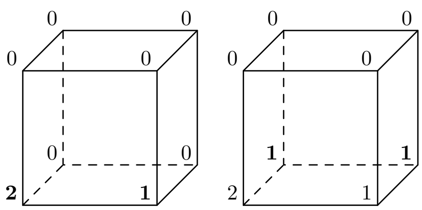

Alas, neither configuration wins the game.

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. That number is an overestimate because the order of the moves does not affect the final state. I could push that 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.*

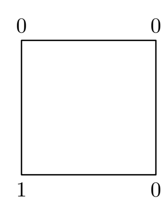

The square is easier to analyze than is the cube, but standards can be lowered still more by analyzing the one-dimensional analog, a line. Having only one edge means that there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that 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

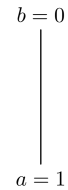

$$a - b \equiv 1 \pmod 3,$$

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,

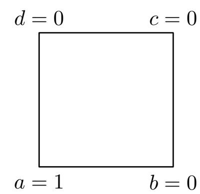

$$a - b + c - d \equiv 1 \pmod 3,$$

so it is impossible to get all vertices to be multiples of 3.

The original, three-dimensional solitaire game is also likely to be impossible to win. 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

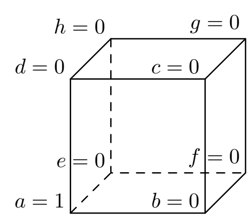

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3,$$

which shows that 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, 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: *When there is change, look for what does not change.*

6.2 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.

6.2.1 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}.

6.2.2 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 $c_{\text{d}}/2$:

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

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

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

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

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

and the total energy to fly is

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

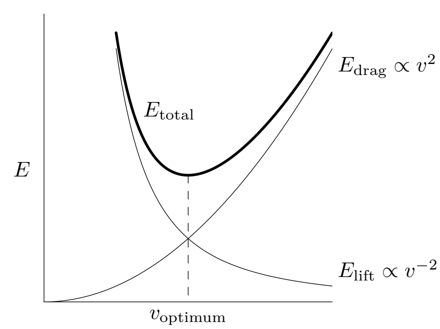

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that 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.

6.2.3 How the maximum range depends on size

Before calculating a range for a particular plane or bird, evaluate the scaling: How does the range depend on the size of the plane? As for the mountain-height analysis (**Section 5.2**), 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 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 km, only slightly longer than the godwit's range.

6.2.4 Explicit computations

To get an explicit range, not only how the range scales with size, 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 $9\,\mathrm{kcal\,g^{-1}}$. In metric units, it is $4 \cdot 10^7\,\mathrm{J\,kg^{-1}}$. Including a typical engine efficiency of one-fourth gives

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

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:

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

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

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

to our definition

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

to get

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

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

Combine the values to find the range:

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

The maximum range of a 747-400 is 13,450 km. The maximum known nonstop flight by a bird – indeed, by any animal – is 11,570 km: A female bar-tailed godwit tracked by satellite migrated between Alaska and New Zealand. The approximate analysis of the range is unreasonably accurate.

Next I estimate the minimum-energy speed and compare it to the cruising speed of a 747. The sum of drag and lift energies is a minimum when the speed is given by

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

The speed is

$$v \sim \left(\frac{Mg}{C^{1/2}\rho L^2}\right)^{1/2}.$$

A fully loaded 747 has $M \sim 4\cdot 10^5$ kg. The drag coefficient is again $C \sim 1/600$, the wingspan is $L \sim 60$ m, and the air density up high is $\rho \sim 0.5\,\mathrm{kg\,m^{-3}}$. So

$$v \sim \left(\frac{4\cdot 10^5\,\mathrm{kg} \times 10\,\mathrm{m\,s^{-2}}}{(1/600)^{1/2} \times 0.5\,\mathrm{kg\,m^{-3}} \times 3.6\cdot 10^3\,\mathrm{m^2}}\right)^{1/2}.$$

Do the arithmetic mentally. The $\sqrt{1/600}$ in the denominator becomes a 25 in the numerator. Combined with the $4\cdot 10^5$, it becomes 10^7. Including the 10 from g, the numerator is 10^8 and the denominator is roughly $2\cdot 10^3$, so

$$v \sim \left(\frac{1}{2}\cdot 10^5\right)^{1/2}\,\mathrm{m\,s^{-1}} = 5^{1/2} \times 100\,\mathrm{m\,s^{-1}} \sim 220\,\mathrm{m\,s^{-1}}.$$

That speed is roughly 500 mph, reasonably close to the 747's maximum speed of 608 mph.

Chapter 7
Dimensions

The next way to eliminate spurious complexity is the method of dimensional analysis or dimensionless groups.

The following mathematical problem shows how dimensionless groups are, like proportional reasoning and conservation laws, a form of symmetry reasoning. The problem is to expand the polynomial $(a+b+c)^3$ into its numerous terms. The school-algebra method is to multiply $a+b+c$ by $a+b+c$ and then multiply the result by $a+b+c$.

The school-algebra method is messy prone to mistakes, but symmetry comes to the rescue. The factor $a+b+c$ is unchanged if a an b, if a and c, or if b and c are swapped. Indeed, any permutation of a, b, and c leaves $a+b+c$ unchanged and leaves $(a+b+c)^3$ unchanged. Therefore, the product can – and should be – built using invariant combinations of a, b, and c: combinations that are unchanged by permuting a, b, and c.

The product contains a^3, but a^3 alone is not invariant to a permutation. The invariant version of a^3 is $a^3+b^3+c^3$, which is one component of the product. The product also contains terms like a^2b, which again is not invariant. The analogous invariant sum comes from adding all possible permutations of a^2b:

$$(a^2b+ab^2)+(a^2c+ac^2)+(b^2c+bc^2).$$

The third type of term is abc.

So the product has the form

$$(a+b+c)^3 = A(a^3+b^3+c^3)+B(a^2b+ab^2+a^2c+ac^2+b^2c+bc^2)+C(abc),$$

where A, B, and C are for-the-moment-unknown constants.

Here is one way to evaluate the constants. Set $a=1$, $b=c=0$. Then the equation reduces to

$$1^3 = A \cdot 1^3,$$

so $A=1$. To get another relation, set $a=b=c=1$. Then, using $A=1$, the equation reduces to

$$27 = 3+6B+C.$$

To find C, notice that there are six ways to get an abc factor. So $C = 6$ and then $B = 3$.

Thus

$$(a + b + c)^3 = (a^3 + b^3 + c^3) + 3(a^2b + ab^2 + a^2c + ac^2 + b^2c + bc^2) + 6abc.$$

This symmetry solution has several merits. First, it is less prone to mistakes than is multiplying by brute force. Second, it produces the answer in a meaningful, low-entropy form. The chunks in the solution – the terms $a^3 + b^3 + c^3$ and $a^2b + ab^2 + a^2c + ac^2 + b^2c + bc^2$ and abc – each obey the symmetry that nothing important changes if you permute a, b, and c. Rather than using a brute-force method and then doing hard work to turn the solution into a meaningful form, use symmetry reasoning: Whenever possible, work with quantities that obey the symmetries of the problem.

This chapter shows how this idea leads naturally to dimensionless groups, the fundamental idea of dimensional analysis.

7.1 Power of multinationals

The first example shows what happens when people take no notice of dimensions.

Critics of globalization often make this argument:

> 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. [Source: 'Impunity for Multinationals', *ATTAC*, 11 Sept 2002, `[url:nigeria-argument]`, retrieved 11 Sept 2006]

Before reading further, try to find the most egregious fault in the comparison between Exxon and Nigeria. It's a competitive field, but one fault stands out.

The comparison between Exxon and Nigeria has many problems. First, the comparison exaggerates Exxon's power by using its worldwide assets (net worth) rather than its assets only in Nigeria. On the other hand, Exxon can use its full international power when negotiating with Nigeria, so perhaps the worldwide assets are a fair basis for comparison.

A more serious, and less debatable, problem is the comparison with GDP, or gross domestic product. To see the problem, look at the ingredients in how GDP is usually measured: as dollars per year. The $99 billion for Nigeria's GDP is shorthand for $99 billion per year. A year is an astronomical time, and its use in an economic measurement is arbitrary. Economic flows, which are a social phenomenon, should not care about how long the earth requires to travel around the sun. Suppose instead that the decade was the chosen unit of time in measuring the GDP. Then Nigeria's GDP would be roughly $1 trillion per decade (assuming that the $99 billion per year value held steady) and would be reported as $1 trillion. Now Nigeria towers over the puny Exxon whose assets are a mere one-tenth of this figure.

To produce the opposite conclusion, just measure GDP in units of dollars per week: Nigeria's GDP becomes \$2 billion per week. Now puny Nigeria stands helpless before the might of Exxon, 50-fold larger than Nigeria. Either conclusion about the relative powers can be produced merely by changing the units. This arbitrariness indicates that the comparison is bogus.

The flaw in the comparison is the theme of this chapter. Assets, or net worth, are an amount of money – money is its dimensions – and are typically measured in units of dollars. GDP is defined as the total goods and services sold in one year. It is a rate and has dimensions of money per time; its typical units are dollars per year. Comparing assets to GDP means comparing money to money per time. *Because the dimensions of these two quantities are not the same, the comparison is nonsense!* A similarly flawed comparison is to compare length per time (speed) with length. Listen how ridiculous it sounds: 'I walk 1.5 meters per second, much smaller than the Empire State building in New York, which is 300 meters high.' To produce the opposite conclusion, measure time in hours: 'I walk 5000 meters per hour, much larger than the Empire State building at only 300 meters.' Nonsense all around!

This example illustrates several ideas:

- *Dimensions versus units.* Dimensions are general and generic, such as money per time or length per time. Units are the instantiation of dimensions in a system of measurement. The most complete system of measurement is the System International (SI), where the unit of mass is the kilogram, the unit of time the second, and the unit of length the meter. Other examples of units are dollars per year or kilometers per year.
- *Necessary condition for a valid comparison.* In a valid comparison, the dimensions of the compared objects be identical. Do not compare apples to oranges (except in questions of taste, like 'I prefer apples to oranges.')
- *Rubbish abounds.* There's lots of rubbish out there, so keep your eyes open for it!
- *Bad argument, fine conclusion.* I agree with the conclusion of the article, that large oil companies exert massive power over poor countries. However, as a physicist I am embarrassed by the reasoning. This example teaches me a valuable lesson about theorems and proofs: judge the proof not just the theorem. Even if you disagree with the conclusion, remember the general lesson that a correct conclusion does not validate a dubious argument.

7.2 Pyramid volume

The last example showed the value of dimensions in economics. The next example shows that dimensions are also useful in mathematics. What is the volume of this square-based pyramid? Here are several choices:

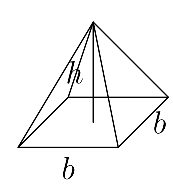

1. $\frac{1}{3}bh$
2. $b^3 + h^2$

3. b^4/h
4. bh^2

Let's take the choices in turn. The first choice, $bh/3$, has dimensions of area rather than volume. So it cannot be right. The second choice, $b^3 + h^2$, begins with a volume in the b^3 term but falls apart with the h^2, which has dimensions of area. Since it adds an area to a volume – the crime of dimension mixing – it cannot be right. The third choice, b^4/h, has dimensions of volume, so it might be correct. It even increases as b increases, which is a good sign. However, the volume should increase as h increases – a proportional-reasoning argument – whereas this choice indicates that the volume decreases as h increases! So it cannot be right.

The final choice, bh^2, has correct dimensions and increases as h or b increases. Does it increase by the right amounts? Imagine drilling into the pyramid from the top and dividing it into thin cores or volume elements. If the height of the pyramid doubles, then each vertical volume element doubles in volume; so the volume of the pyramid should double. In symbols, $V \propto h$. But bh^2 quadruples when h doubles, so that choice cannot be right.

The requirement that $V \propto h$ together with the requirement that V have dimensions of length cubed means that the missing item in $V \propto h$ is an area. The only way to make an area from b is to make b^2 perhaps times a dimensionless constant. So

$$V \sim hb^2.$$

The missing dimensionless constant is hidden in the twiddle $\sim$ sign. Alternatively, the ratio V/hb^2 is dimensionless.

This method of deducing the volume requires remembering hardly any arbitrary data. It requires these ingredients:

1. Using vertical volume elements to find out that $V \propto h$.
2. Using dimensions along with $V \propto h$ to show that $V \sim hb^2$.
3. Remembering the correct dimensionless constant.

The first two steps are logic and do not require arbitrary data. Instead they use reasoning methods that you use elsewhere (so there's no marginal cost to remember them). The third step requires seemingly arbitrary data. However, in **Chapter 8** on special cases, I'll show you how to determine the constant elegantly without even needing an integral.

Then the volume requires no memory. Arbitrary data is, by definition, impossible to compress. Dimensions, and more generally our techniques for handling complexity, are a form of data compression or entropy reduction [5]. One way to look at learning is as data compression. So dimensions, and our other techniques, enhance learning.

There's an old saying: Tell the truth; there's less to remember. The similar moral here is: *Use dimensions (and proportional reasoning); there's less to remember!*

7.3 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 true statement about these variables looks like

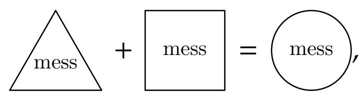

where the various messes mean 'a horrible combination of E, s, v, ρ, L, and m.

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

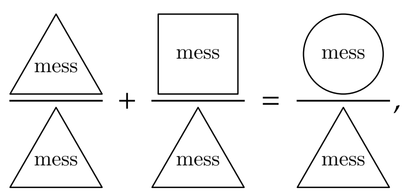

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

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

Therefore, *any true 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. Here is a familiar 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. **Section 5.1** gave a proportional-reasoning analysis of this system. 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.

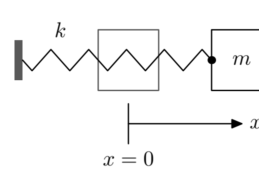

Var	*Dim*	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

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. The analysis in **Section 5.1** gave an approximate argument for why the period should be independent of the amplitude. So that approximate argument turns out to be an exact argument.
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 [6, 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!

7.4 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, 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* [8].) 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.

7.4.1 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}	*surfacetension*		

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

Then

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

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

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

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

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

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

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, so the dimensional analysis has given the exact answer.

7.4.2 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

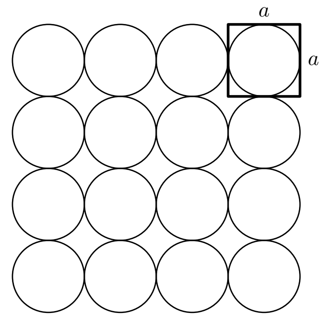

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

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

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

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 $10\,\text{g}\,\text{cm}^{-3}$. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is $3\,\text{g}\,\text{cm}^{-3}$, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid $3000\,\text{kg}\,\text{m}^{-3}$. 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 $\text{g}\,\text{cm}^{-3}$. 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

7.4.3 Physical interpretation

The previous method, dimensional analysis, is mostly mathematical. As a second computation of a_0, we show you a method that is mostly physics. Besides checking the Bohr radius, it provides a physical interpretation of it. 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**:

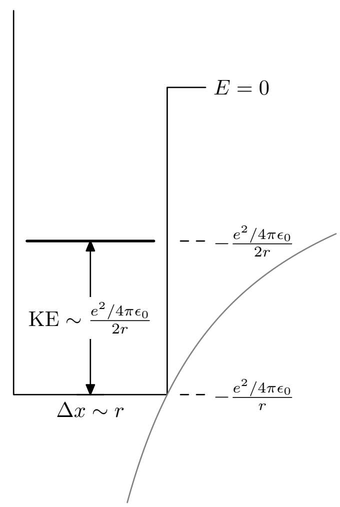

$$\Delta x \Delta p \sim \hbar,$$

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 idea recurs in the book.

The potential energy is the classical expression

$$E_{\text{Potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

The total energy is the combination

$$E = E_{\text{Potential}} + E_{\text{Kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

The two energies compete. At small r, kinetic energy wins, because of the $1/r^2$; at large r, potential energy wins, because it goes to zero less rapidly. Is there a minimum combined energy at some intermediate value of r? There has to be. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope crosses between positive and negative. The energy is a a minimum there. The location would be easy to estimate if the energy were written in dimensionless form. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is worth learning in this example.

In constructing the dimensionless group containing a_0, we constructed another length:

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

To scale any length – to make it dimensionless – divide it by l. So in the total energy the scaled radius

$$\bar{r} \equiv \frac{r}{l}.$$

The other unknown in the total energy is the energy itself. To make it dimensionless, a reasonable energy scale to use is $e^2/4\pi\epsilon_0 l$ by defining scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the dimensionless length and energy, the total energy

$$E = E_{\text{Potential}} + E_{\text{Kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

becomes

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants are placed into the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

Simple calculus: minimizing scaled energy $\bar{E}$ versus scaled bond length $\bar{r}$. The scaled energy is the sum of potential and kinetic energy. The shape of this energy illustrates Feynman's explanation of the atomic hypothesis. At a 'little distance apart' – for large $\bar{r}$ – the curve slopes upward; to lower their energy, the proton and electron prefer to move closer, and the resulting force is attractive. 'Upon being squeezed into one another' – for small $\bar{r}$ – the potential rapidly increases, so the force between the particles is repulsive. Somewhere between the small and large regions of $\bar{r}$, the force is zero.

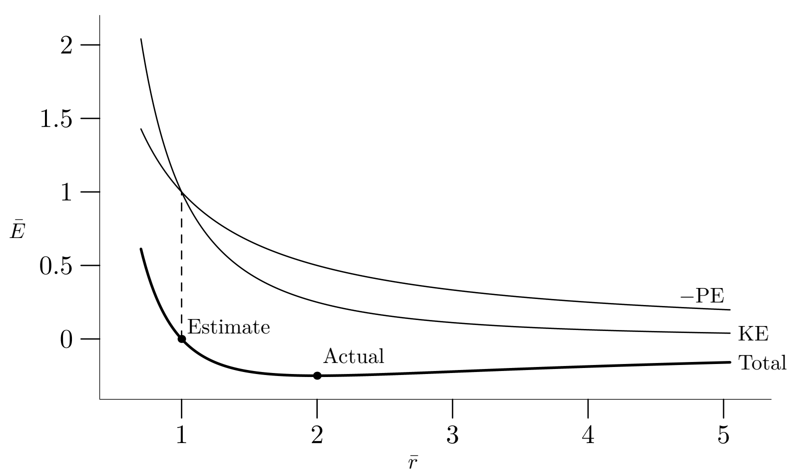

Calculus (differentiation) locates this minimum-energy $\bar{r}$ at $\bar{r}_{\text{min}} = 2$. An alternative method is **cheap minimization**: When two terms compete, the minimum occurs when the terms are roughly equal. This method of minimization is familiar from **Section 6.2.2**.

Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{\text{min}} \sim 1$. This result gives a scaled length. In actual units, it is

$$r_{\text{min}} = l\bar{r}_{\text{min}} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius computed using dimensional analysis. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization!

Here is how to justify cheap minimization. Consider a reasonable general form for E:

$$E(r) = \frac{A}{r^n} - \frac{B}{r^m}.$$

This form captures the important feature of the combined energy

$$E = E_{\text{Potential}} + E_{\text{Kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2},$$

that two terms represent competing physical effects. Mathematically, that physical fact is shown by the opposite signs.

To find the minimum, solve $E'(r_{\text{min}}) = 0$ or

$$-n\frac{A}{r_{\text{min}}^{n+1}} + m\frac{B}{r_{\text{min}}^{m+1}} = 0.$$

The solution is

$$\frac{A}{r_{\text{min}}^n} = \frac{n}{m}\frac{B}{r_{\text{min}}^m} \qquad \text{(calculus)}.$$

This method minimizes the combined energy by equating the two terms A/r^n and B/r^m:

$$\frac{A}{r_{\text{min}}^n} = \frac{B}{r_{\text{min}}^m}.$$

This approximation lacks the n/m factor in the exact result. The ratio of the two estimates for r_{min} is

$$\frac{\text{approximate estimate}}{\text{calculus estimate}} \sim \left(\frac{n}{m}\right)^{1/(m-n)},$$

which is smaller than 1 unless $n = m$, when there is no maximum or minimum. So the approximate method underestimates the location of minima and maxima.

To judge the method in practice, apply it to a typical example: the potential between nonpolar atoms or molecules, such as between helium, xenon, or methane. This potential is well approximated by the so-called Lennard–Jones potential where $m = 6$ and $n = 12$:

$$U(r) \sim \frac{A}{r^{12}} - \frac{B}{r^6}.$$

The approximate result underestimates r_{min} by a factor of

$$\left(\frac{12}{6}\right)^{1/6} \sim 1.15.$$

An error of 15 percent is often small compared to the other inaccuracies in an approximate computation, so this method of approximate minimization is a valuable time-saver.

Now return to the original problem: determining the Bohr radius. The approximate minimization predicts the correct value. Even if the method were not so charmed, there is no point in doing a proper, calculus minimization. The calculus method is *too accurate* given the inaccuracies in the rest of the derivation.

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 instead of 50 years, but such a replacement would be overengineering. To improve a bicycle, put effort into improving or doing without the welds.

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.

This method of equating competing terms is **balancing**. We balanced the kinetic energy against the potential energy by assuming that they are roughly the same size. The consequence is that

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

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. ‘Nearly 1’ is also called **of order unity**. This rough equality occurs in many examples, and you often get a reasonable answer 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).

7.5 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. I can calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

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

To compute the Christoffel symbols $\Gamma^{\beta}_{\mu\nu}$ requires solving for the metric tensor $g_{\mu\nu}$, which requires solving the curvature equations $R_{\mu\nu} = 0$.

The curvature equations are a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous – that's one problem – but worse, they are not linear. So the standard trick, which is to guess a type of solution and form new solutions by combining the basic types, does not work. You can spend a decade learning advanced mathematics to solve the equations exactly. Or you can accept the great principle of analysis: When the going gets tough, lower your standards. If I sacrifice accuracy, I can explain light bending in less 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.

The following sections do each step.

7.5.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. An unlabeled diagram prods me into thinking of labels, many of which are parameters of the problem:

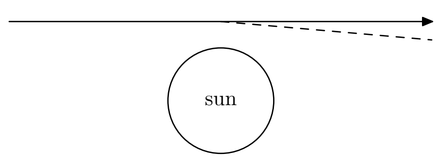

Here are reasons to include various parameters:

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 same diagram with important parameters labeled:

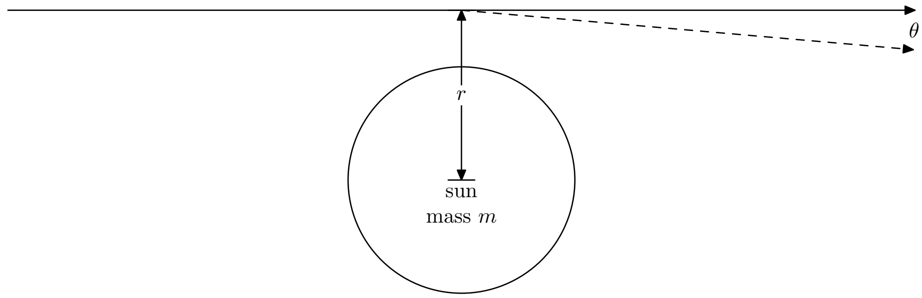

Here is a table of the parameters and their dimensions:

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, as you might suspect, L, M, and T represent the dimensions of length, mass, and time, respectively.

7.5.2 Dimensionless groups

What are the dimensionless groups? 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.

I want a second dimensionless group because otherwise my analysis seems like nonsense. Any physical solution can be written in dimensionless form; this idea is the foundation of dimensional analysis. With only one dimensionless group, θ, I have to conclude that θ depends on no variables at all:

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

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. My physical picture might be confused, but it's not so confused that neither variable is relevant. So I need to make another dimensionless group on which θ can depend. Therefore, I return to Step 1: Finding parameters.

The list so far lacks a crucial parameter.

What physics have I 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 my parameters distinguish between light and rocks. Therefore I should include c, the speed of light. It introduces the fact that I'm 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	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	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 I already know I need a combination such as Gm to cancel out mass. Time also appears in only two parameters: G and c. To cancel out time, I need to form Gm/c^2. This combination has one length in it, so a dimensionless group is Gm/rc^2.

7.5.3 Drawing conclusions

The most general relation between the two dimensionless groups is

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

Dimensional analysis cannot tell me the correct function f.

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 try symmetry: Imagine a world where gravity is repulsive or, equivalently, the gravitational constant is negative. Then the angle should also be negative, so f should be an odd function. This symmetry argument eliminates choices like $f(Gm/rc^2) \sim (Gm/rc^2)^2$.

The simplest guess is that f is the identity function. Then the bending angle is

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

There is likely a dimensionless constant in f:

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

or

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

are also possible. This freedom means

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

7.5.4 Comparison with exact calculations

Different theories of gravity give the same result

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

the only variation is in the value for the missing dimensionless constant. Here are those values from exact calculation:

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

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just that. 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 twice the Newtonian value because light moves at the speed of light. The extra bending is a consequence of Einstein's theory of special relativity putting space and time on the same level. The theory of general relativity then formulates gravity in terms of the curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; general relativity itself also calculates the space curvature. Since most objects move much slower than the speed of light, meaning that they travel much farther in time than in space, they feel mostly the time curvature. The Newtonian analysis is fine for those objects. Since light moves at the speed of light, it sees equal amounts of space and time curvature, so it bends twice as far as the Newtonian theory would predict.

7.5.5 Numbers!

At the surface of the Earth, the 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). So 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 I quote 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. 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 lenses made is 6 m. So there is no chance of detecting an 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, tried to measure exactly this effect.

For many years Einstein believed that his theory of gravity would predict the Newtonian value, which turns out to be 0.87 arcseconds for light just grazing 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, which predicted 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, so 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, which could measure small deflections accurately. I leave you with this puzzle: If the accuracy of a telescope is λ/d, how could radio telescopes be more accurate than optical ones, since radio waves have a longer wavelength than light has?!

7.6 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	$\mathrm{MLT^{-2}}$	force
m	M	mass
a	$\mathrm{LT^{-2}}$	acceleration

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

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

And that statement is the Buckingham Pi theorem [9].

Part 3

Discarding Information

Chapter 8
Special cases

8.1 Pyramid volume

I have been promising to explain the factor of one-third in the volume of a pyramid:

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

Although the method of special cases mostly cannot explain a dimensionless constant, the volume of a pyramid provides a rare counterexample.

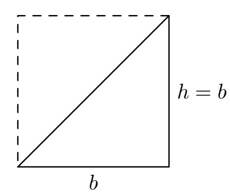

I first explain the key idea in fewer dimensions. So, instead of immediately explaining the one-third in the volume of a pyramid, which is a difficult three-dimensional problem, first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h: The area is $A \sim bh$. What is the constant? Choose a convenient triangle, perhaps a 45-degree right triangle where $h = b$. Two of those 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 $A = bh/2$.

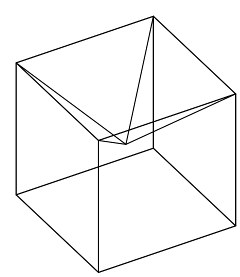

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? Only the aspect ratio h/b matters in the following discussion. So choose b conveniently, and then choose h to make a pyramid with the clever aspect ratio. The goal shape is suggested by the square pyramid base. Another solid with the same base is a cube.

Perhaps several pyramids can combine 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 the 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:

$$V = \frac{1}{3}hb^2 = \frac{4}{3}.$$

8.2 Mechanics

8.2.1 Atwood machine

The next problem illustrates dimensional analysis and special cases in a physical problem. Many of the ideas and methods from the geometry example transfer to this problem, and it introduces more methods and ways of reasoning.

The problem is a staple of first-year physics: 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, educated guessing, and a feel for functions.

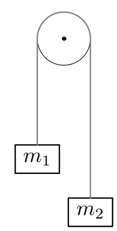

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is 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 special 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 $a = 0$.

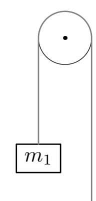

Here is a plot of our knowledge of f:

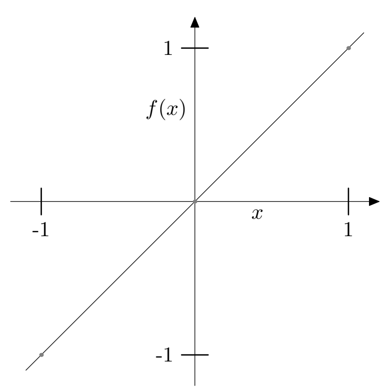

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

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

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

8.3 Drag

Pendulum motion is not a horrible enough problem to show the full benefit of dimensional analysis. Instead try fluid mechanics – a subject notorious for its mathematical and physical complexity; Chandrasekhar's books [10, 11] or the classic textbook of Lamb [12] show that the mathematics is not for the faint of heart.

The next examples illustrate two extremes of fluid flow: oozing and turbulent. An example of oozing flow is ions transporting charge in seawater (**Section 8.3.6**). An example of turbulent flow is a raindrop falling from the sky after condensing out of a cloud (**Section 8.3.7**).

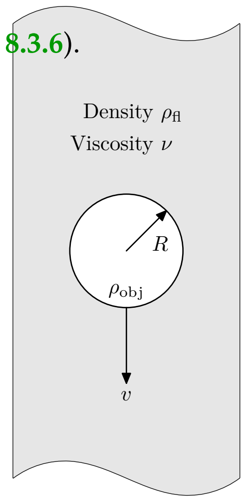

To find the terminal velocity, solve the partial-differential equations of fluid mechanics for the incompressible flow of a Newtonian fluid:

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v}, \qquad \text{(3 eqns)}$$

$$\boldsymbol{\nabla}\cdot\mathbf{v} = 0. \qquad \text{(1 eqn)}$$

Here $\mathbf{v}$ is the fluid velocity, ρ is the fluid density, ν is the kinematic viscosity, and p is the pressure. The first equation is a vector shorthand for three equations, so the full system is four equations.

All the equations are partial-differential equations and three are nonlinear. Worse, they are coupled: Quantities appear in more than one equation. So we have to solve a system of coupled, nonlinear, partial-differential equations. This solution must satisfy boundary conditions imposed by the marble or raindrop. As the object moves, the boundary conditions change. So until you know how the object moves, you do not know the boundary conditions. Until you know the boundary conditions, you cannot find the motion of the fluid or of the object. This coupling between the boundary conditions and solution compounds the difficulty of the problem. It requires that you solve the equations and the boundary conditions together. If you ever get there, then you take the limit $t \to \infty$ to find the terminal velocity.

Sleep easy! I wrote out the Navier–Stokes equations only to scare you into using dimensional analysis and special-cases reasoning. The approximate approach is easier than solving nonlinear partial-differential equations.

8.3.1 Naive dimensional analysis

To use dimensional analysis, follow the usual steps: Choose relevant variables, form dimensionless groups from them, and solve for the terminal velocity. In choosing quantities, do not forget to include the variable for which you are solving, which here is v. To decide on the other quantities, split them into three categories (divide and conquer):

1. characteristics of the fluid,
2. characteristics of the object, and
3. characteristics of whatever makes the object fall.

The last category is the easiest to think about, so deal with it first. Gravity makes the object fall, so g is on the list.

Consider next the characteristics of the object. Its velocity, as the quantity for which we are solving, is already on the list. Its mass m affects the terminal velocity: A feather falls more slowly than a rock does. Its radius r probably affects the terminal velocity. Instead of listing r and m together, remix them and use r and ρ_{obj}. The two alternatives r and m or r and ρ_{obj} provide the same information as long as the object is uniform: You can compute ρ_{obj} from m and r and can compute m from ρ_{obj} and r.

Choose the preferable pair by looking ahead in the derivation. The relevant properties of the fluid include its density ρ_{fl}. If the list also includes ρ_{obj}, then the results might contain pleasing dimensionless ratios such as $\rho_{\text{obj}}/\rho_{\text{fl}}$ (a dimensionless group!). The ratio $\rho_{\text{obj}}/\rho_{\text{fl}}$ has a more obvious physical interpretation than a combination such as $m/\rho_{\text{fl}}r^3$, which, except for a dimensionless constant, is more obscurely the ratio of object and fluid densities. So choose ρ_{obj} and r over m and r.

Scaling arguments also favor the pair ρ_{obj} and r. In a scaling argument you imagine varying, say, a size. Size, like heat, is an extensive quantity: a quantity related to amount of stuff. When you vary the size, you want as few other variables as possible to change so that those changes do not obscure the effect of changing size. Therefore, whenever possible replace extensive quantities with **intensive quantities** like temperature or density. The pair m and r contains two extensive quantities, whereas the preferable pair ρ_{obj} and r contains only one extensive quantity.

Now consider properties of the fluid. Its density ρ_{fl} affects the terminal velocity. Perhaps its viscosity is also relevant. Viscosity measures the tendency of a fluid to reduce velocity differences in the flow. You can observe an analog of viscosity in traffic flow on a multilane highway. If one lane moves much faster than another, drivers switch from the slower to the faster lane, eventually slowing down the faster lane. Local decisions of the drivers reduce the velocity gradient. Similarly, molecular motion (in a gas) or collisions (in a fluid) transports speed (really, momentum) from fast- to slow-flowing regions. This transport reduces the velocity difference between the regions. Oozier (more viscous) fluids probably produce more drag than thin fluids do. So viscosity belongs on the list of relevant variables.

Fluid mechanicians have defined two viscosities: dynamic viscosity η and kinematic viscosity ν. [Sadly, we could not use the mellifluous term *fluid mechanics* to signify a host of physicists agonizing over the equations of fluid mechanics; it would not distinguish the toilers from their toil.] The two viscosities are related by $\eta = \rho_{\text{fl}}\nu$. *Life in Moving Fluids* [13, pp. 23–25] discusses the two types of viscosity in detail. For the analysis of drag force, you need to know only that viscous forces are proportional to viscosity. Which viscosity should we use? Dynamic viscosity hides ρ_{fl} inside the product $\nu\rho_{\text{fl}}$; a ratio of ρ_{obj} and η then looks less dimensionless than it is because ρ_{obj}'s partner ρ_{fl} is buried inside η. Therefore the kinematic viscosity ν usually gives the more insightful results. Summarizing the discussion, the table lists the variables by category.

Var	*Dim*	What
ν	$\mathrm{L}^2\mathrm{T}^{-1}$	kinematic viscosity
ρ_{fl}	ML^{-3}	fluid density
r	L	object radius
v	LT^{-1}	terminal velocity
ρ_{obj}	ML^{-3}	object density
g	LT^{-2}	gravity

The next step is to find dimensionless groups. The Buckingham Pi theorem (**Section 7.6**) says that the six variables and three independent dimensions result in three dimensionless groups.

Before finding the groups, consider the consequences of three groups. Three?! Three dimensionless groups produce this form for the terminal velocity v:

$$\text{group with } v = f(\text{other group 1}, \text{other group 2}).$$

To deduce the properties of f requires physics knowledge. However, studying a two-variable function is onerous. A function of one variable is represented by a curve and can be graphed on a sheet of paper. A function of two variables is represented by a surface. For a complete picture it needs three-dimensional paper (do you have any?); or you can graph many slices of it on regular two-dimensional paper. Neither choice is appealing. This brute-force approach to the terminal velocity produces too many dimensionless groups.

If you simplify only after you reach the complicated form

$$\text{group with } v = f(\text{other group 1}, \text{other group 2}),$$

you carry baggage that you eventually discard. When going on holiday to the Caribbean, why pack skis that you never use but just cart around everywhere? Instead, at the beginning of the analysis, incorporate the physics knowledge. That way you simplify the remainder of the derivation. To follow this strategy of packing light – of packing only what you need – consider the physics of terminal velocity in order to make simplifications now.

8.3.2 Simpler approach

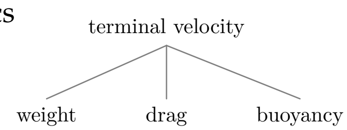

The adjective *terminal* in the phrase 'terminal velocity' hints at the physics that determines the velocity. Here 'terminal' is used in its sense of final, as in after an infinite time. It indicates that the velocity has become constant, which happens only when no net force acts on the marble. This line of thought suggests that we imagine the forces acting on the object: gravity, buoyancy, and drag. The terminal velocity is velocity at which the drag, gravitational, and buoyant forces combine to make zero net force. Divide-and-conquer reasoning splits the terminal-velocity problem into three simpler problems.

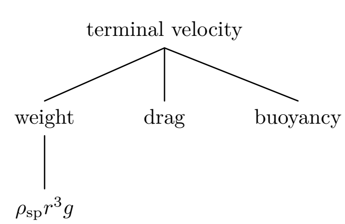

The gravitational force, also known as the weight, is mg. Instead of m we use $(4\pi/3)\rho_{\text{obj}}r^3$ – for the same reasons that we listed ρ_{obj} instead of m in the table of variables – and happily ignore the factor of $4\pi/3$. With those choices, the weight is

$$F_{\text{g}} \sim \rho_{\text{obj}} r^3 g.$$

The figure shows the roadmap updated with this information.

The remaining pieces are drag and buoyancy. Buoyancy is easier, so do it first (the principle of maximal laziness). It is an upward force that results because gravity affects the pressure in a fluid. The pressure increases according to $p = p_0 + \rho_{\text{fl}} g h$, where h is the depth and p_0 is the pressure at zero depth (which can be taken to be at any level in the fluid). The pressure difference between the top and bottom of the object, which are separated by a distance $\sim r$, is $\Delta p \sim \rho_{\text{fl}} g r$. Pressure is force per area, and the pressure difference acts over an area $A \sim r^2$. Therefore the buoyant force created by the pressure difference is

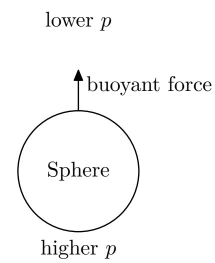

$$F_{\text{b}} \sim A\Delta p \sim \rho_{\text{fl}} r^3 g.$$

As a check on this result, Archimedes's principle says that the buoyant force is the 'weight of fluid displaced'. This weight is

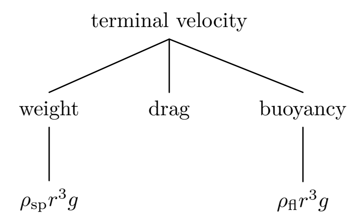

$$\overbrace{\rho_{\text{fl}} \underbrace{\frac{4\pi}{3}\pi r^3}_{\textit{volume}}}^{\textit{mass}} g.$$

Except for the factor of $4\pi/3$, it matches the buoyant force so Archimedes's principle confirms our estimate for F_{b}. That result updates the roadmap. The main unexplored branch is the drag force, which we solve using dimensional analysis.

8.3.3 Dimensional analysis for the drag force

The weight and buoyancy were solvable without dimensional analysis, but we still need to use dimensional analysis to find the drag force. The purpose of breaking the problem into parts was to simplify this dimensional analysis relative to the brute-force approach in **Section 8.3.1**. Let's see how the list of variables changes when computing the drag force rather than the terminal velocity. The drag force F_{d} has to join the list: not a promising beginning when trying to eliminate variables. Worse, the terminal velocity v remains on the list, even though we are no longer computing it, because the drag force depends on the velocity of the object.

However, all is not lost. The drag force has no idea what is inside the sphere. Picture the fluid as a huge computer that implements the laws of fluid dynamics. From the viewpoint of this computer, the parameters v and r are the only relevant attributes of a moving sphere. What lies underneath the surface does not affect the fluid flow: Drag is only skin deep. The computer can determine the flow (if it has tremendous processing power) without knowing the sphere's density ρ_{obj}, which means it vanishes from the list. Progress!

Var	*Dim*	What
F_d	MLT^{-2}	drag force
ν	L^2T^{-1}	kinematic viscosity
ρ_{fl}	ML^{-3}	fluid density
r	L	object radius
v	LT^{-1}	terminal velocity

Now consider the characteristics of the fluid. The fluid supercomputer still needs the density and viscosity of the fluid to determine how the pieces of fluid move in response to the object's motion. So ρ_{fl} and ν remain on the list. What about gravity? It causes the object to fall, so it is responsible for the terminal velocity v. However, the fluid supercomputer does not care how the object acquired this velocity; it cares only what the velocity is. So g vanishes from the list. The updated tabled shows the new, shorter list.

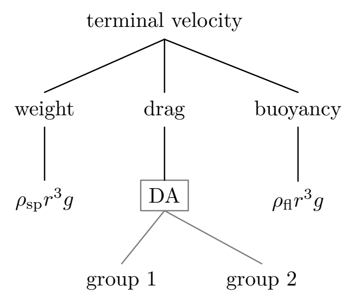

The five variables in the list are composed of three basic dimensions. From the Buckingham Pi theorem (**Section 7.6**), we expect two dimensionless groups. We find one group by dividing and conquering. The list already includes a velocity (the terminal velocity). If we can concoct another quantity V with dimensions of velocity, then v/V is a dimensionless group. The viscosity ν is almost a velocity. It contains one more power of length than velocity does. Dividing by r eliminates the extra length: $V \equiv \nu/r$. A dimensionless group is then

$$G_1 \equiv \frac{v}{V} = \frac{vr}{\nu}.$$

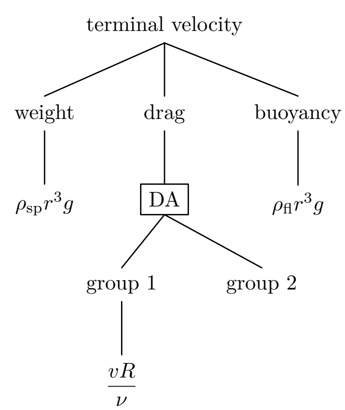

Our knowledge, including this group, is shown in the figure. This group is so important that it has a name, the **Reynolds number**, which is abbreviated *Re*. It is important because it is a *dimensionless* measure of flow speed. The velocity, because it contains dimensions, cannot distinguish fast from slow flows. For example, $1000\,\mathrm{m\,s^{-1}}$ is slow for a planet, whose speeds are typically tens of kilometers per second, but fast for a pedestrian. When you hear that a quantity is small, fast, large, expensive, or almost any adjective, your first reaction should be to ask, 'compared to what?' Such a comparison suggests dividing v by another velocity; then we get a dimensionless quantity that is proportional to v. The result of this division is the Reynolds number.

Low values of *Re* indicate slow, viscous flow (cold honey oozing out of a jar). High values indicate turbulent flow (a jet flying at 600 mph). The excellent *Life in Moving Fluids* [13] discusses many more dimensionless ratios that arise in fluid mechanics.

The Reynolds number looks lonely in the map. To give it company, find a second dimensionless group. The drag force is absent from the first group so it must live in the second; otherwise we cannot solve for the drag force.

Instead of dreaming up the dimensionless group in one lucky guess, we construct it in steps (divide-and-conquer reasoning). Examine the variables in the table, dimension by dimension. Only two (F_d and ρ_{fl}) contain mass, so both or neither appear in the group. Because F_d has to appear, ρ_{fl} must also appear. Each variable contains a first power of mass, so the group contains the ratio F_d/ρ_{fl}. A simple choice is

$$G_2 \propto \frac{F_d}{\rho_{fl}}.$$

The dimensions of F_d/ρ_{fl} are L^4T^{-2}, which is the square of L^2T^{-1}. Fortune smiles on us, for L^2T^{-1} are the dimensions of ν. So

$$\frac{F_d}{\rho_{fl}\nu^2}$$

is a dimensionless group.

This choice, although valid, has a defect: It contains ν, which already belongs to the first group (the Reynolds number). Of all the variables in the problem, ν is the one most likely to be found irrelevant based on a physical argument (as will happen in **Section 8.3.7**, when we specialize to high-speed flow. If ν appears in two groups, eliminating it requires recombining the two groups into one that does not contain ν. However, if ν appears in only one group, then eliminating it is simple: eliminate that group. Simpler mathematics – eliminating a group rather than remixing two groups to get one group – requires simpler physical reasoning. Therefore, isolate ν in one group if possible.

To remove ν from the proposed group $F_d/\rho_{fl}\nu^2$ notice that the product of two dimensionless groups is also dimensionless. The first group contains ν^{-1} and the proposed group contains ν^{-2}, so the ratio

$$\frac{\text{group proposed}}{(\text{first group})^2} = \frac{F_d}{\rho_{fl} r^2 v^2}$$

is not only dimensionless but it also does not contain ν. So the analysis will be easy to modify when we try to eliminate ν. With this revised second group, our knowledge is now shown in this figure:

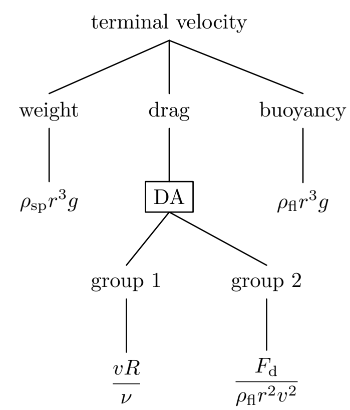

This group, unlike the the proposal $F_d/\rho_{fl}\nu^2$, has a plausible physical interpretation. Imagine that the sphere travels a distance l, and use l to multiply the group by unity:

$$\underbrace{\frac{F_d}{\rho_{fl} r^2 v^2}}_{\text{group 1}} \times \underbrace{\frac{l}{l}}_{1} = \frac{F_d l}{\rho_{fl} l r^2 v^2}.$$

The numerator is the work done against the drag force over the distance l. The denominator is also an energy. To interpret it, examine its parts (divide and conquer). The product lr^2 is, except for a dimensionless constant, the volume of fluid swept out by the object. So $\rho_{fl} l r^2$ is, except for a constant, the mass of fluid shoved aside by the object. The object moves fluid with a velocity comparable to v, so it imparts to the fluid a kinetic energy

$$E_K \sim \rho_{fl} l r^2 v^2.$$

Thus the ratio, and hence the group, has the following interpretation:

$$\frac{\text{work done against drag}}{\text{kinetic energy imparted to the fluid}}.$$

In highly dissipative flows, when energy is burned directly up by viscosity, the numerator is much larger than the denominator, so this ratio (which will turn out to measure drag) is much greater than 1. In highly streamlined flows (a jet wing), the the work done against drag is small because the fluid returns most of the imparted kinetic energy to the object. So in the ratio, the numerator will be small compared to the denominator.

To solve for F_d, which is contained in G_2, use the form $G_2 = f(G_1)$, which becomes

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The drag force is then

$$F_d = \rho_{fl} r^2 v^2 f\left(\frac{vr}{\nu}\right).$$

The function f is a dimensionless function: Its argument is dimensionless and it returns a dimensionless number. It is also a universal function. The same f applies to spheres of any size, in a fluid of any viscosity or density! Although f depends on r, ρ_{fl}, ν, and v, it depends on them only through one combination, the Reynolds number. A function of one variable is easier to study than is a function of four variables:

> 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.
>
> —Harold Jeffreys [6, p. 82]

Dimensional analysis cannot tell us the form of f. To learn its form, we specialize to two special cases:

1. viscous, low-speed flow ($Re \ll 1$), the subject of **Section 8.3.4**; and
2. turbulent, high-speed flow ($Re \gg 1$), the subject of **Section 8.3.7**.

8.3.4 Viscous limit

As an example of the low-speed limit, consider a marble falling in vegetable oil or glycerin. You may wonder how often marbles fall in oil, and why we bother with this example. The short answer to the first question is 'not often'. However, the same physics that determines the fall of marbles in oil also determines, for example, the behavior of fog droplets in air, of bacteria swimming in water [14], or of oil drops in the Millikan oil-drop experiment. The marble problem not only illustrates the physical principles, but also we can check our results with a home experiment.

In slow, viscous flows, the drag force comes directly from – surprise! – viscous forces. These forces are proportional to viscosity because viscosity is the constant of proportionality in the definition of the viscous force. Therefore

$$F_d \propto \nu.$$

The viscosity appears exactly once in the drag result, repeated here:

$$F_d = \rho_{fl} r^2 v^2 f\left(\frac{vr}{\nu}\right).$$

To flip ν into the numerator and make $F_d \propto \nu$, the function f must have the form $f(x) \sim 1/x$. With this $f(x)$ the result is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanician must do a messy and difficult calculation. Her burden is light now that we have worked out the solution except for this one constant. The British mathematician Stokes, the first to derive its value, found that

$$F_d = 6\pi \rho_{fl} \nu v r.$$

In honor of Stokes, this result is called Stokes drag.

Let's sanity check the result. Large or fast marbles should feel a lot of drag, so r and v should be in the numerator. Viscous fluids should produce a lot of drag, so ν should be the numerator. The proposed drag force passes these tests. The correct location of the density – in the numerator or denominator – is hard to judge.

You can make an educated judgment by studying the Navier–Stokes equations. In those equations, when v is 'small' (small compared to what?) then the $(\mathbf{v}\cdot\nabla)\mathbf{v}$ term, which contains two powers of v, becomes tiny compared to the viscous term $\nu\nabla^2\mathbf{v}$, which contains only one power of v. The second-order term arises from the inertia of the fluid, so this term's being small says that the oozing marble does not experience inertial effects. So perhaps ρ_{fl}, which represents the inertia of the fluid, should not appear in the Stokes drag. On the other hand, viscous forces are proportional to the *dynamic* viscosity $\eta = \rho_{fl}\nu$, so ρ_{fl} should appear even if inertia is unimportant. The Stokes drag passes this test. Using the dynamic instead of kinematic viscosity, the Stokes drag is

$$F_d = 6\pi\eta v r,$$

often a convenient form because many tables list η rather than ν.

This factor of 6π comes from doing honest calculations. Here, it comes from solving the Navier–Stokes equations. In this book we wish to teach you how not to suffer, so we do not solve such equations. We usually quote the factor from honest calculation to show you how accurate (or sloppy) the approximations are. The factor is often near unity, although not in this case where it is roughly 20! In fancy talk, it is usually 'of order unity'. Such a number suits our neural hardware: It is easy to remember and to use. Knowing the approximate derivation and remembering this one number, you reconstruct the exact result without solving difficult equations.

Now use the Stokes drag to estimate the terminal velocity in the special case of low Reynolds number.

8.3.5 Terminal velocity for low Reynolds number

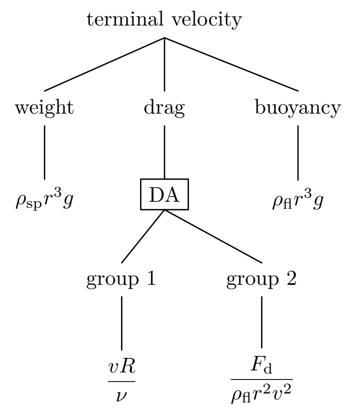

Having assembled all the pieces in the roadmap, we now return to the original problem of finding the terminal velocity. Since no net force acts on the marble (the definition of terminal velocity), the drag force plus the buoyant force equals the weight:

$$\underbrace{\nu\rho_{fl}vr}_{F_d} + \underbrace{\rho_{fl}gr^3}_{F_b} \sim \underbrace{\rho_{obj}gr^3}_{F_g}.$$

After rearranging:

$$\nu\rho_{fl}vr \sim (\rho_{obj} - \rho_{fl})gr^3.$$

The terminal velocity is then

$$v \sim \frac{gr^2}{\nu}\left(\frac{\rho_{obj}}{\rho_{fl}} - 1\right).$$

In terms of the dynamic viscosity η, it is

$$v \sim \frac{gr^2}{\eta}(\rho_{obj} - \rho_{fl}).$$

This version, instead of having the dimensionless factor $\rho_{obj}/\rho_{fl} - 1$ that appears in the version with kinematic viscosity, has a dimensional $\rho_{obj} - \rho_{fl}$ factor. Although it is less aesthetic, it is often more convenient because tables often list dynamic viscosity η rather than kinematic viscosity ν.

We can increase our confidence in this expression by checking whether the correct variables are upstairs (a picturesque way to say 'in the numerator') and downstairs (in the denominator). Denser marbles should fall faster than less dense marbles, so ρ_{obj} should live upstairs. Gravity accelerates marbles, so g should live upstairs. Viscosity slows marbles, so ν should live downstairs. The terminal velocity passes these tests. We therefore have more confidence in our result, although the tests did not check the location of r or any exponents: For example, should ν appear as ν^2? Who knows, but if viscosity matters, it mostly appears as a square root or as a first power.

To check r, imagine a large marble. It will experience a lot of drag and fall slowly, so r should appear downstairs. However, large marbles are also heavy and fall rapidly, which suggests that r should appear upstairs. Which effect wins is not obvious, although after you have experience with these problems, you can make an educated guess: weight scales as r^3, a rapidly rising function r, whereas drag is probably proportional to a lower power of r. Weight usually wins such contents, as it does here, leaving r upstairs. So the terminal velocity also passes the r test.

Let's look at the dimensionless ratio in parentheses: $\rho_{obj}/\rho_{fl} - 1$. Without buoyancy the -1 disappears, and the terminal velocity would be

$$v \propto g\,\frac{\rho_{\text{obj}}}{\rho_{\text{fl}}}.$$

We retain the g in the proportionality for the following reason: The true solution returns if we replace g by an effective gravity g' where

$$g' \equiv g\left(1 - \frac{\rho_{\text{fl}}}{\rho_{\text{obj}}}\right).$$

So, one way to incorporate the effect of the buoyant force is to solve the problem without buoyancy but with the reduced g.

Check this replacement in two limiting cases: $\rho_{\text{fl}} = 0$ and $\rho_{\text{fl}} = \rho_{\text{obj}}$. When $\rho_{\text{obj}} = \rho_{\text{fl}}$ gravity vanishes: People, whose density is close to the density of water, barely float in swimming pools. Then g' should be zero. When $\rho_{\text{fl}} = 0$, buoyancy vanishes and gravity retains its full effect. So g' should equal g. The effective gravity definition satisfies both tests. Between these two limits, the effective g should vary linearly with ρ_{fl} because buoyancy and weight superpose linearly in their effect on the object. The effective g passes this test as well.

Another test is to imagine $\rho_{\text{fl}} > \rho_{\text{obj}}$. Then the relation correctly predicts that g' is negative: helium balloons rise. This alternative to using buoyancy explicitly is often useful. If, for example you forget to include buoyancy (which happened in the first draft of this chapter), you can correct the results later by replacing g with the g'.

If we carry forward the constants of proportionality, starting with the magic 6π in the Stokes drag and including the $4\pi/3$ that belongs in the weight, we find

$$v \sim \frac{2}{9}\,\frac{gr^2}{\nu}\left(\frac{\rho_{\text{obj}}}{\rho_{\text{fl}}} - 1\right).$$

8.3.6 Conductivity of seawater

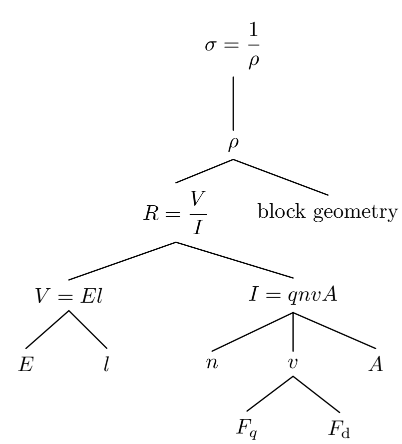

As an application of Stokes drag and a rare example of a realistic situation with low Reynolds numbers, let's estimate the electrical conductivity of seawater. Solving this problem is hopeless without breaking it into pieces. Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the convention that overloads the density symbol with yet another meaning.) Resistivity, as its name suggests, is related to resistance R. Why have both ρ and R? Resistance is a useful measure for a particular wire, but not for wires in general because it depends on the diameter and cross-sectional area of the wire. It is not an intensive quantity. Before examining the relationship between resistivity and resistance, let's finish sketching the solution tree, leaving ρ as depending on R plus geometry. We can find R by placing a voltage V across a block of seawater and measuring the current I; then $R = V/I$.

To find V or I we need a physical model. First, why does seawater conduct at all? Conduction requires the transport of charge, which is produced by an electric field. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt can 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 the ion concentration, and v is its terminal speed.

To understand, and be able to rederive this formula, first check its dimensions. Current is charge per time. Is the right side also charge per time? Yes: q takes care of the charge; and vA has dimensions of L^3T^{-1} so nvA, which has dimensions of T^{-1}, takes care of the time.

As a second check, watch a cross-section of the block for a time Δt. How much charge flows in that time? The charges move at speed v, so all charges in block of width $v\Delta t$ and area A cross the cross-section. This 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 flow, so the current is $I = Q/\Delta t = qnvA$. The terminal speed v depends on the applied force F_q and on the drag force F_d, just as for the falling marble but with an electrical force instead of a gravitational force. The result of this subdividing is the preceding map.

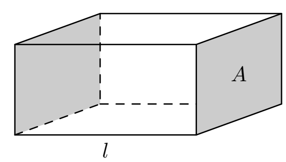

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

$$\rho = \frac{RA}{l}.$$

To find v we follow the same procedure as for the marble. 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

$$F_q = \frac{qV}{l},$$

an expression in terms only of known quantities. The drag is Stokes drag. Equating this drag 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.

Only the number density n remains unknown. We estimate it after getting a symbolic result for σ, which you can do by climbing up the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2 nAV}{6\pi\eta l r}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta l r}{q^2 nA}.$$

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:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2 n}.$$

The expression simplifies as we rise up the tree: The geometric parameters l and A have also vanished, which is also encouraging: The purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier, so

$$\sigma = \frac{1}{\rho} \sim \frac{e^2 n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. This value I estimate with 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 a teaspoon of salt tastes very salty, 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 $5\,\mathrm{m}\ell$. For those who live in metric countries, a teaspoon is an archaic measure used in Britain and especially the United States, which has no nearby metric country to which it pays attention. 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 ~ 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

$$\begin{aligned} n &\sim \frac{1}{30} \times \underbrace{1\,\mathrm{g\,cm^{-3}}}_{\rho_{\text{water}}} \times \frac{2\,\text{charges}}{\text{molecule}} \times \frac{6\cdot 10^{23}\,\text{molecules mole}^{-1}}{60\,\mathrm{g\,mole^{-1}}} \\ &\sim 7\cdot 10^{20}\,\text{charges}\,\mathrm{cm^{-3}}. \end{aligned}$$

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

$$\eta = \rho_{\text{water}}\nu \sim 10^3\,\text{kg}\,\text{m}^{-3} \times 10^{-6}\,\text{m}^2\,\text{s}^{-1} = 10^{-3}\,\text{kg}\,\text{m}^{-1}\,\text{s}^{-1}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units – also known as God's units – than they are 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 ghastly μ_0 or $4\pi\epsilon_0$, for example – than it is in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations, 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 an educated 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.

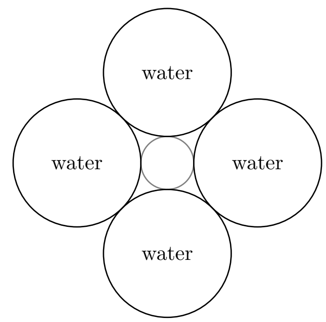

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6\cdot 10^{-19}\,\text{C})^2}^{e^2} \times \overbrace{7\cdot 10^{26}\,\text{m}^{-3}}^{n}}{\underbrace{6\times 3}_{6\pi} \times \underbrace{10^{-3}\,\text{kg}\,\text{m}^{-1}\,\text{s}^{-1}}_{\eta} \times \underbrace{2\cdot 10^{-10}\,\text{m}}_{r}}.$$

You can do the computation mentally: *Take out the big part*, apply the *principle of maximal laziness*, and *divide and conquer* by first counting the powers of ten (shown in red) and then worrying about the small factors. Then *divide and conquer* again by counting the top and bottom contributions separately. The top contributes -12 powers of ten: -38 from e^2 and $+26$ from n. The bottom contributes -13 powers of ten: -3 from η and -10 from r. The division produces one power of ten.

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

$$\sigma \sim 5\,\Omega^{-1}\,\text{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\text{m}^{-1}$, 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 $2\,\mathrm{\AA}$, 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 more accurate results than it ought to.

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 illustrating Stokes drag, is to have courage. Approximate first and ask questions later. Maybe the approximations are correct 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.

8.3.7 Turbulent limit

We now compute drag in the other flow extreme: high-speed, or turbulent, flow. The example will be to compute the terminal speed of a raindrop. These results apply to most flows. For example, when a child rises from a chair, the airflow around her is high-speed flow, as you can check by computing the Reynolds number. Say that the child is $0.2\,\mathrm{m}$ wide, and that she rises with velocity $0.5\,\mathrm{m\,s^{-1}}$. Then

$$Re \sim \frac{vr}{\nu_{\mathrm{air}}} \sim \frac{0.5\,\mathrm{m\,s^{-1}} \times 0.2\,\mathrm{m}}{2\cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}} \sim 5000.$$

Here viscosity of air is closer to

$$\nu_{\mathrm{air}} \approx 1.5\cdot 10^{-5}\,\mathrm{m^2\,s^{-1}},$$

than to $2\cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}$, but $2\cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}$ easily combines with the $0.2\,\mathrm{m}$ in the numerator to allows us to do the calculation mentally. Using either value for the viscosity, the Reynolds number is much larger than unity, so the flow is turbulent. Larger objects, such as planes, trains, and automobiles, create turbulence even when they travel even more slowly than the child. In short, most fluid flow around us is turbulent flow.

To begin the analysis, we assume that a raindrop is a sphere. It is a convenient lie that allows us to reuse the general results of **Section 8.3.3** and specialize to high-speed flow. At high speeds (more precisely, at high Reynolds number) the flow is turbulent. Viscosity – which affects only slow flows but does not directly influence the shearing and whirling of turbulent flows – becomes irrelevant. Let's see how much we can understand about turbulent drag knowing only that turbulent drag is nearly independent of viscosity.

Turbulence is perhaps the main unsolved problem in classical physics. However, you can still understand a lot about drag using dimensional analysis plus a bit of physical reasoning; we do not need a full understanding of turbulence. The world is messy: Do not wait for a full understanding before you analyze or estimate.

In the roadmap for low Reynolds number, the viscosity appears only in the first group. Because turbulent drag is independent of the viscosity, the viscosity disappears from the results and therefore so does that group. This argument is glib. More precisely, remove ν from the list of variables and search again for dimensionless groups. The remaining four variables, shown in the table, result in one dimensionless group, which is the second group from the old roadmap.

Var	*Dim*	What
F_d	MLT^{-2}	drag force
ρ_{fl}	ML^{-3}	fluid density
r	L	object radius
v	LT^{-1}	terminal velocity

So the Reynolds number, which was the first group, has disappeared from the analysis. But why is drag at high speeds independent of Reynolds number? Equivalently, why can we remove ν from the list of variables and still get the correct form for the drag force? The answer is not obvious. The explanation of the Reynolds number as a ratio of two speeds v and V provides a partial answer. A natural length in this problem is r; we can use r to transform v and V into times:

$$\tau_v \equiv \frac{r}{v},$$
$$\tau_V \equiv \frac{r}{V} \sim \frac{r^2}{\nu}.$$

Note that $Re \equiv \tau_V/\tau_v$. The quantity τ_v is the time that fluid takes to travel around the sphere (apart from constants). Kinematic viscosity is ν/ρ, but its most important interpretation is as the diffusion coefficient for momentum. So the time for momentum to diffuse a distance x is

$$\tau \sim \frac{x^2}{\nu}.$$

This result depends on the mathematics of random walks; you can increase your confidence in it here, without understanding the theory of random walks, by checking that it has valid dimensions. And it has: Each side is a time.

So τ_V is the time that momentum takes to diffuse around an object of size r, such as the falling sphere in this problem. If $\tau_V \ll \tau_v$ – in which case $Re \ll 1$ – then momentum diffuses before fluid travels around the sphere. Momentum diffusion equalizes velocities, if it has time, which it does have in this low-Reynolds-number limit. Momentum diffusion therefore prevents flow at the front from being radically different from the flow at the back, and thereby squelches any turbulence. In the other limit, when $\tau_V \gg \tau_v$ or $Re \gg 1$ – momentum diffusion is outraced by fluid flow, so the fluid is free to shred itself into a turbulent mess. Once the viscosity is low enough to allow turbulence, its value does not affect the drag, which is why we can ignore it for $Re \gg 1$. Here $Re \gg 1$ means 'large enough so that turbulence sets in', which happens around $Re \sim 1000$. A more complete story, which we discuss as part of boundary layers in **Section 9.4**, slightly corrects this approximation. However, it is close enough for our purposes here.

Here the important point is that the viscosity vanishes from the analysis and so does group 1. Once it disappears, the dimensionless group that remains is

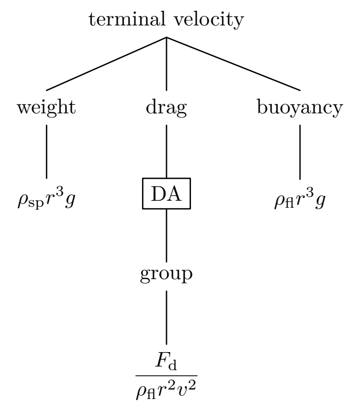

$$G_2 = \frac{F_d}{\rho_{fl} r^2 v^2}.$$

Because it is the only group, the solution is

$$G_2 = \text{dimensionless constant},$$

or

$$F_d \sim \rho_{fl} r^2 v^2.$$

This drag is for a sphere. What about other shapes, which are characterized by more parameters than a sphere is? So that the drag force generalizes to more complex shapes, we express it using the cross-sectional area of the object. Here $A = \pi r^2$, so

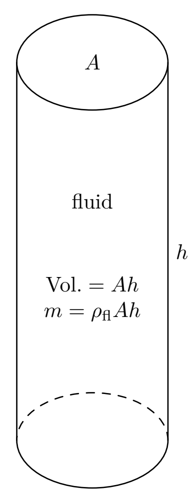

$$F_d \sim \rho_{fl} A v^2.$$

This conventional choice has a physical basis. As an object moves, the mass of fluid that it displaces is proportional to its cross-sectional area:

$$m_{fl} = \rho_{fl} A h.$$

The fluid is given a speed comparable to v, so the fluid's kinetic energy is

$$E_K \sim \frac{1}{2} m_{fl} v^2 \sim \frac{1}{2} \rho_{fl} A h v^2.$$

If all this kinetic energy is dissipated by drag, then the drag force is E_K/h or

$$F_d \sim \frac{1}{2} \rho_{fl} A v^2.$$

In this form with the factor of 1/2, the constant of proportionality is the **drag coefficient** c_d.

Object	c_d
Sphere	0.5
Cylinder	1.0
Flat plate	2.0
Car	0.4

Like its close cousin f from the dimensionless drag force, the drag coefficient is a dimensionless measure of the drag force. It depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number). The drag coefficient, being proportional to the function $f(Re)$ in the general solution, also depends on the Reynolds number. However, using the reasoning that the flow at high Reynolds number is independent of viscosity, the drag coefficient should also be independent of Reynolds number. Using the drag coefficient instead of f (which implies using cross-sectional area instead of r^2), the turbulent drag force becomes

$$F_d = \frac{1}{2}c_d\rho_{fl}v^2A.$$

So we have an expression for the turbulent drag force. The weight and buoyant forces are the same as in the viscous limit. So we just need to redo the analysis of the viscous limit but with the new drag force. Because the weight and buoyant forces contain r^3, we return to using r^2 instead of A in the drag force. With these results, the terminal velocity v is given by

$$\underbrace{\rho_{fl}r^2v^2}_{F_d} \sim \underbrace{g(\rho_{obj}-\rho_{fl})r^3}_{F_g-F_b},$$

so

$$v \sim \sqrt{gr\left(\frac{\rho_{obj}}{\rho_{fl}}-1\right)}.$$

Pause to sanity check this result: Are the right variables upstairs and downstairs? We consider each variable in turn.

- ρ_{fl}: The terminal velocity is smaller in a denser fluid (try running in a swimming pool), so ρ_{fl} should be in the denominator.
- g: Imagine a person falling on a planet that has a gravitational force stronger than that of the earth. Gravity partially determines atmospheric pressure and density. Holding the atmospheric density constant while increasing gravity might be impossible in real life, but we can do it easily in a thought experiment. The drag force then does not depend on g, so gravity increases the terminal speed without opposition from the drag force: g should be upstairs.
- ρ_{obj}: Imagine a raindrop made of (very) heavy water. Relative to a standard raindrop, the gravitational force increases while the drag force remains constant, as shown using the fluid-is-a-computer argument in ??sec:drag-force-DA. So ρ_{obj} should be upstairs.
- r: To determine where the radius lives requires a more subtle argument. Increasing r increases both the gravitational and drag forces. The gravitational force increases as r^3

whereas the drag force increases only as r^2. So, for larger raindrops, their greater weight increases v more than their greater drag decreases v. Therefore r should be live upstairs.

- ν: At high Reynolds number viscosity does not affect drag, at least not in our approximation. So ν should not appear anywhere.

The terminal velocity passes all tests.

Now we can compute the terminal velocity. The splash spots on the sidewalk made by raindrops in a recent rain have $r \sim 0.3\,\text{cm}$. Since rain is water, its density is $\rho_{\text{obj}} \sim 1\,\text{g}\,\text{cm}^{-3}$. The density of air is $\rho_{\text{fl}} \sim 1\,\text{kg}\,\text{m}^{-3}$, so $\rho_{\text{fl}} \ll \rho_{\text{obj}}$: Buoyancy is therefore not an important effect, and we can replace $\rho_{\text{obj}}/\rho_{\text{fl}} - 1$ by $\rho_{\text{obj}}/\rho_{\text{fl}}$. With this simplification and the estimated numbers, the terminal velocity is:

$$v \sim \left(\underbrace{1000\,\text{cm}\,\text{s}^{-2}}_{g} \times \underbrace{0.3\,\text{cm}}_{r} \times \frac{\overbrace{1\,\text{g}\,\text{cm}^{-3}}^{\rho_{\text{obj}}}}{\underbrace{10^{-3}\,\text{g}\,\text{cm}^{-3}}_{\rho_{\text{fl}}}} \right)^{1/2} \sim 5\,\text{m}\,\text{s}^{-1},$$

or 10 mph.

This calculation assumed that $Re \gg 1$. Check that assumption! You need not calculate Re from scratch; rather, scale it relative to a previous results. As we worked out earlier, a child ($r \sim 0.2\,\text{m}$) rising from her chair ($v \sim 0.5\,\text{m}\,\text{s}^{-1}$) creates a turbulent flow with $Re \sim 5000$. The flow created by the raindrop is faster by a factor of 10, but the raindrop is smaller by a factor of roughly 100. Scaling the Reynolds number for the child gives

$$Re \sim \underbrace{Re_{\text{child}}}_{5000} \times \underbrace{\left(\frac{v_{\text{drop}}}{v_{\text{child}}}\right)}_{10} \times \underbrace{\left(\frac{r_{\text{drop}}}{r_{\text{child}}}\right)}_{0.01} \sim 500.$$

This Reynolds number is also much larger than 1, so the flow produced by the raindrop is turbulent, which vindicates the initial assumption.

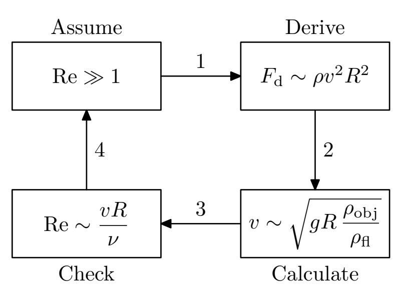

Now that we have found the terminal velocity, let's extract the pattern of the solution. The order that we followed was *assume*, *derive*, *calculate*, then *check*. This order is more fruitful than is the simpler order of *derive* then *calculate*. Without knowing whether the flow is fast or slow, we cannot derive a closed-form expression for F_d; such a derivation is probably beyond present understanding of fluids and turbulence. Blocked by this mathematical Everest, we would remain trapped in the *derive* box. We would never determine F_d, so we would never realize that the Reynolds number is large (the *assume* box); however, only this assumption makes it possible to eliminate ν and thereby to estimate F_d. The moral: Assume early and often!

8.3.8 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. Compared to the turbulent form, the viscous form lacks one power of r and one power of v but has an extra power of ν. A combination of variables with a similar property is the Reynolds number rv/ν. So multiply the viscous drag by a useful form of unity:

$$F_\mathrm{d} = \underbrace{\left(\frac{rv/\nu}{Re}\right)}_{1} \times \underbrace{6\pi\rho_\mathrm{fl}v\rho_\mathrm{fl}\nu r}_{F_\mathrm{d}} = \frac{1}{Re}\,6\pi\rho_\mathrm{fl}v^2r^2 \qquad \text{(viscous)}.$$

This form, except for the 6π and the r^2, resembles the turbulent drag Fortunately $A = \pi r^2$ so

$$F_\mathrm{d} = \frac{6}{Re}\,\rho_\mathrm{fl}v^2A \qquad \text{(viscous)},$$

With

$$c_\mathrm{d} = \frac{12}{Re} \qquad \text{(viscous)},$$

the turbulent drag and this rewritten viscous drag for a sphere have the same form:

$$F_\mathrm{d} = \frac{1}{2}\rho_\mathrm{fl}Av^2 \times \begin{cases} \dfrac{12}{Re} & (Re \ll 1), \\ 0.5 & (Re \gg 1). \end{cases}$$

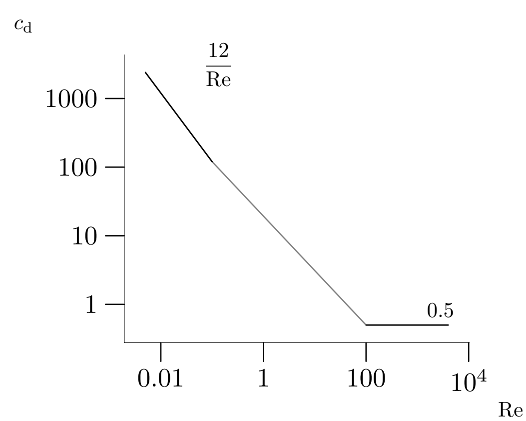

At high Reynolds number the drag coefficient remains constant. For a sphere, that constant is $c_\mathrm{d} \sim 1/2$. If the low-Reynolds-number approximation for c_d is valid at sufficiently high Reynolds numbers, then c_d would cross 1/2 near $Re \sim 24$, where presumably the high-Reynolds-number approximation takes over. The crossing point is a reasonable estimate for the transition between low- and high-speed flow. Experiment or massive simulation are the only ways to get a more accurate result. Experimental data place the crossover near $Re \sim 5$, at which point $c_\mathrm{d} \sim 2$. Why can't you calculate this value analytically? If a dimensionless variable, such as the Reynolds number, is close to unity, calculations become difficult. Approximations that depend on a quantity being either huge or tiny are no longer valid. When all terms in an equation are roughly of the same magnitude, you cannot get rid of any term without making large errors. To get results in these situations, you have to do honest work: You must do experiments or solve the Navier–Stokes equations numerically.

Chapter 9
Discretization

9.1 Diaper usage

9.2 Pendulum period

9.3 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, and your leg twitches. How long does it take to get there? 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. In a physical random walk, for example a gas molecule moving and colliding, the walker moves a variable distance and can move in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details.

Simplify by discarding the generality. The generality arises from the continuous degrees of freedom: the direction is continuous and the distance between collisions is continuous. So, discretize the direction and the distance: Assume that the particle travels a fixed distance between collisions and that it can move only along the coordinate axes. Furthermore, analyze the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this discretized, 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. Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. Because the random walk is unbiased – because moving in each direction is equally likely – the expected position remains constant:

$$\langle x_n \rangle = \langle x_{n-1} \rangle \,.$$

So $\langle x \rangle$, the so-called first moment of the position, is an invariant. However, it is not a fascinating invariant because it does not tell us much that we do not already understand intuitively.

Given that the first moment is not interesting, try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. This analysis is easiest in special cases. Suppose that after a while wandering, 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 in 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,

$$\langle x_n^2 \rangle = n.$$

Since each step takes a constant time, in this discretized analysis, the conclusion is that

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

The result that $\langle x^2 \rangle$ is proportional to time applied to the one-dimensional random walk. And it works for any dimension. 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:

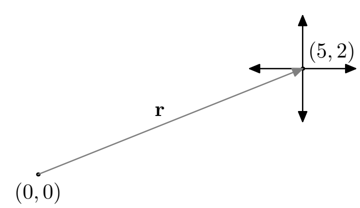

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6, 2)$ and $(4, 2)$. The expected squared distance is

$$\langle x^2 \rangle = \frac{1}{2}(40 + 20) = 30,$$

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.

9.3.1 Difference between a random walk and a regular walk

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 also 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, but the ratio $\langle x^2 \rangle / t$ is an invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* often denoted D. It has dimensions of $\mathrm{L}^2\mathrm{T}^{-1}$.

The 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 a random walker should take longer than a regular walker would take to travel the same distance. 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.

9.3.2 Diffusion equation

The discretized model of a random explains where the diffusion equation comes from. Imagine a gas of particles with each particle doing a random walk in one dimension. How does the concentration, or number, change with time?

Slice the one-dimensional world into slices of width Δx, and 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 at x changes from $N(x)$ to

$$\frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)),$$

for a change of

$$\begin{aligned} \Delta N &= \frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)) - N(x) \\ &= \frac{1}{2}(N(x - \Delta x) - 2N(x) + N(x + \Delta x)). \end{aligned}$$

This last relation can be rewritten as

$$\Delta N \sim (N(x + \Delta x) - N(x)) - (N(x) - N(x + \Delta x)),$$

which in terms of derivatives 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 smashes only curvature – roughly speaking, the second derivative – and does not try to change just the gradient. Heat often diffuses by a random walk, either via phonons (in a liquid or solid) or via molecular random walks (in a gas), so if you maintain one end of a bar at T_1 and the other end at T_2, then the bar will eventually linearly interpolate between the two temperatures, as long as heat is fed into the hot end and drawn out of the cold end.

9.3.3 Keeping warm

One consequence of random walks is how to keep warm on a cold day. We need to calculate the flux of heat: the energy flowing per unit area per unit time. We start from the definition of flux and reason physically.

Flux of stuff 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:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

You can check that the dimensions are the same on both sides.

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass, and $\rho c_p T$ is 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:

$$\text{thermal flux} = \rho c_{\text{p}} T \times \frac{\kappa}{l}.$$

The product $\rho c_{\text{p}} \kappa$ occurs so frequently that it is given a name: the thermal conductivity K. And the ratio T/l is a discretized version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

$$F = K\frac{\Delta T}{\Delta x}.$$

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_{\text{p}} \kappa$. Time to put all the pieces together for air:

$$\begin{aligned}\rho &\sim 1\,\text{kg}\,\text{m}^{-3},\\ c_{\text{p}} &\sim 10^3\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1},\\ \kappa &\sim 1.5\cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1},\end{aligned}$$

where we are guessing that $\kappa = \nu$, since both are diffusion constants. Then

$$K = \rho c_{\text{p}} \kappa \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}.$$

Now we can estimate the heat loss outside on a cold day. Let's say that your skin is at 30 °C and the air outside is 0 °C, so $\Delta T = 30\,\text{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1} \times \frac{30\,\text{K}}{2\cdot 10^{-3}\,\text{m}} \sim 300\,\text{W}\,\text{m}^{-2}.$$

Damn, I wanted a power not a power per area. Oh, flux is power per area, so all is well. I just need to multiply by my surface area. I'm roughly 2 m tall (approximately!) and 0.5 m wide, so my front and back each have area $1\,\text{m}^2$. Then

$$P \sim FA = 300\,\text{W}\,\text{m}^{-2} \times 2\,\text{m}^2 = 600\,\text{W}.$$

No wonder it feels so cold! Just sitting around, your body generates 100 W (the basal metabolic rate). So, with 600 W escaping, you lose far more heat more than you generate. After long enough, your core body temperature drops. Chemical reactions in your body slow down, because all reactions go slower at lower temperature, and because enzymes lose their optimized shape. 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 waste heat, is vigorous enough exercise to keep you 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. Since $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, when wearing the fleece,

$$P \sim 60\,\text{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that you do not feel too cold. Indeed, when wearing a thick fleece, you feel most cold in your hands and face. Those regions are exposed to the air, and are protected by only a thin layer of still air. Because a small Δx means a large heat flux, the moral is: Listen to your parents, bundle up!

9.4 Boundary layers

Chapter 10
Springs

Everything is a spring! The main example in this chapter is waves, which illustrate springs, discretization, and special cases – a fitting, unified way to end the book.

10.1 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. The themes of section are: *Springs are everywhere* and *Consider limiting cases.*

10.1.1 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

$$\bar{\lambda} \equiv \frac{\lambda}{2\pi}.$$

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

$$\omega = ck,$$

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**.

10.1.2 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:

$$f(x, t) = \cos(kx - \omega t),$$

where f is its amplitude.

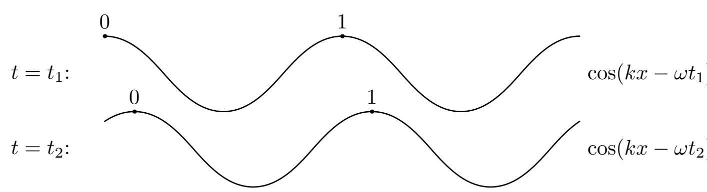

Phase velocity is an easier idea than group velocity so, as an example of divide-and-conquer reasoning and of maximal laziness, study it first. The phase, which is the argument of the cosine, is $kx - \omega t$. A crest occurs when the phase is zero. In the top wave, a crest occurs when $x = \omega t_1/k$. In the bottom wave, a crest occurs when $x = \omega t_2/k$. The difference

$$\frac{\omega}{k}(t_2 - t_1)$$

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

$$v_{\text{ph}} = \frac{\text{distance crest shifted}}{\text{time taken}} = \frac{\omega}{k}.$$

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:

$$f_1 = \cos(kx - \omega t),$$
$$f_2 = \cos((k + \Delta k)x - (\omega + \Delta\omega)t),$$

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.

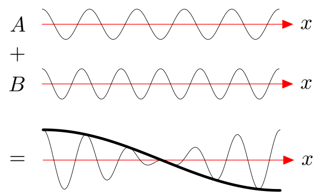

The point of the figure is that two cosines with almost the same spatial frequency add to produce an envelope (thick line). The envelope itself looks like a cosine. After waiting a while, each wave changes because of the ωt or $(\omega + \Delta\omega)t$ terms in their phases. So the sum and its envelope change to this:

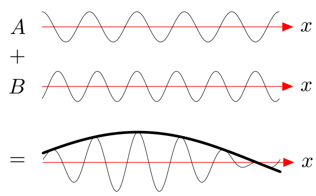

The envelope moves, like the crests and troughs of any wave. It is a wave, so it has a phase velocity, which motivates the following definition:

Group velocity is the phase velocity of the envelope.

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:

$$\cos(A+B) = \underbrace{2\cos\left(\frac{B-A}{2}\right)}_{\textit{envelope}} \times \underbrace{\cos\left(\frac{A+B}{2}\right)}_{\textit{inner}}.$$

If $A \approx B$, then $A - B \approx 0$, which makes the envelope vary slowly. Apply the identity to the sum:

$$f_1 + f_2 = \cos(\underbrace{kx - \omega t}_{A}) + \cos(\underbrace{(k+\Delta k)x - (\omega + \Delta\omega)t}_{B}).$$

Then the envelope contains

$$\cos\left(\frac{B-A}{2}\right) = \cos\left(\frac{x\Delta k - t\Delta\omega}{2}\right).$$

The envelope represents a wave with phase

$$\frac{\Delta k}{2}x - \frac{\Delta \omega}{2}t.$$

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 $f_1 + f_2$:

$$v_g = \frac{\text{frequency}}{\text{wavenumber}} = \frac{\Delta\omega/2}{\Delta k/2} = \frac{\Delta\omega}{\Delta k}.$$

In the limit where $\Delta k \to 0$ and $\Delta\omega \to 0$, the group velocity is

$$v_g = \frac{\partial\omega}{\partial k}.$$

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

$$\frac{v_g}{v_{ph}} = n. \qquad \text{(for a power-law relation)}$$

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.

10.1.3 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.

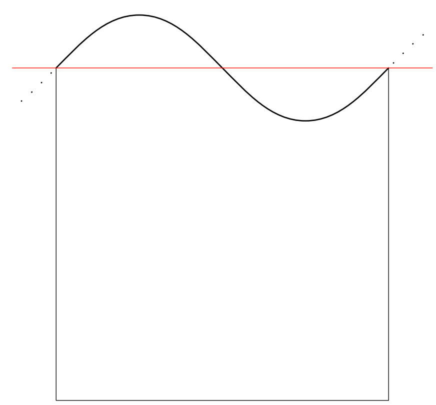

Waves have the necessary reservoirs to act like springs. The surface of water is flat in its lowest-energy state. Deviations, also known as waves, are opposed by a restoring force. Distorting the surface is like stretching a rubber sheet: Surface tension resists the distortion. Distorting the surface 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 [16, 17]. 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,

$$\mathcal{M} \equiv \frac{\text{flow speed}}{\text{sound speed}},$$

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

$$\Pi_1 \equiv kh.$$

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:

$$\Pi_2 = \frac{\omega^2}{gk}.$$

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:

$$\Pi_3 \equiv \frac{F_\gamma}{F_\mathrm{g}} = \frac{\gamma/k}{\rho g/k^3} = \frac{\gamma k^2}{\rho g}.$$

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

$$\Pi_2 = f(\Pi_1, \Pi_3),$$

or

$$\frac{\omega^2}{gk} = f\left(kh, \frac{\gamma k^2}{\rho g}\right).$$

Then

$$\omega^2 = gk \cdot f(kh, \frac{\gamma k^2}{\rho g}).$$

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.

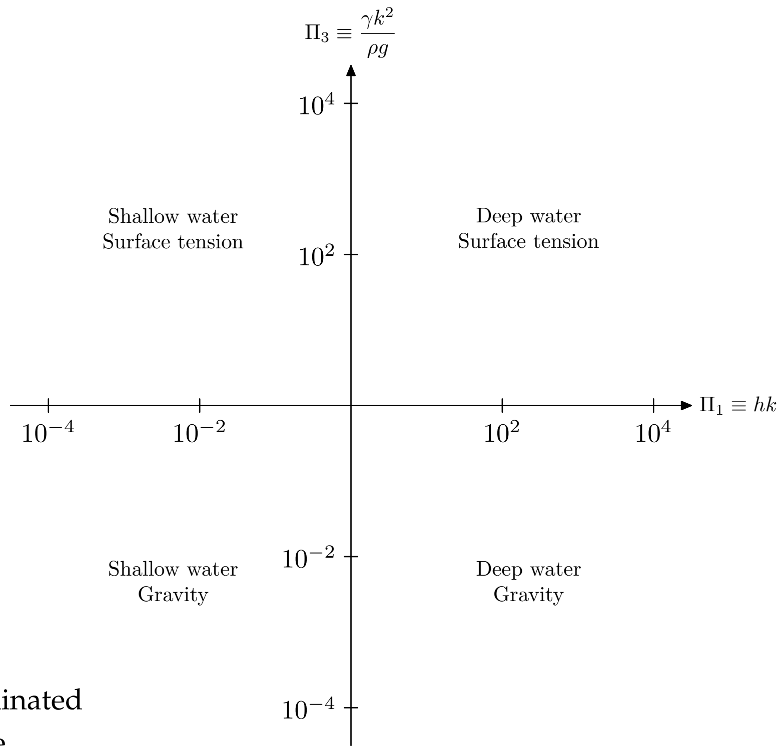

The figure shows how the two groups Π_1 and Π_3 divide the world of waves into four regions. We study each region in turn, and combine the analyses to understand the whole world (of waves). The group Π_1 measures the depth of the water: Are the waves moving on a puddle or an ocean? The group Π_3 measures 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.

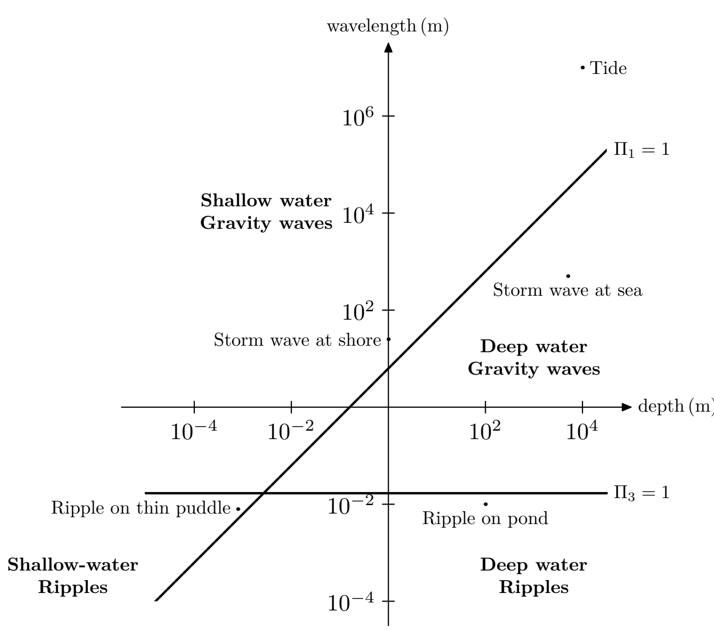

The thick dividing lines are based on the dimensionless groups $\Pi_1 = hk$ and $\Pi_3 = \gamma k^2/\rho g$. Each region contains one or two examples of its kind of waves. With $g = 1000\,\text{cm}\,\text{s}^{-1}$

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.

10.1.4 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:

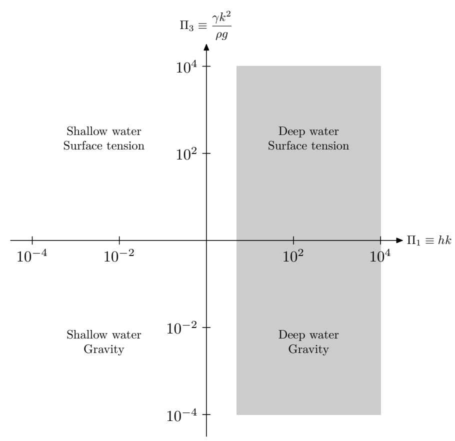

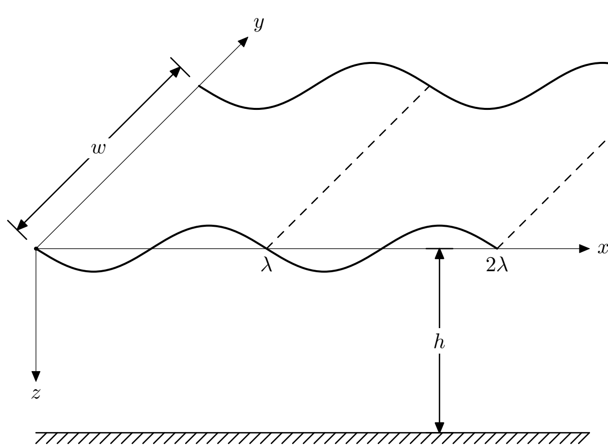

Then, Laplace's equation becomes

$$\frac{\partial^2\phi}{\partial x^2} + \frac{\partial^2\phi}{\partial z^2} = 0,$$

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* [18]; Lamb's *Hydrodynamics* [19] 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

$$\phi \sim Z(z)\sin kx.$$

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*:

$$\frac{\partial^2\phi}{\partial x^2} = -k^2\phi.$$

In order that this ϕ satisfy Laplace's equation, the z-derivative term must contribute $+k^2\phi$. Therefore,

$$\frac{\partial^2\phi}{\partial z^2} = k^2\phi,$$

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 $d \sim 1/k$:

$$\frac{\text{water depth}}{\text{feeler depth}} \sim \frac{h}{1/k} = hk,$$

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

$$\omega^2 = gk\, f_{\text{deep}}\left(\frac{\gamma k^2}{\rho g}\right).$$

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.

10.1.5 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:

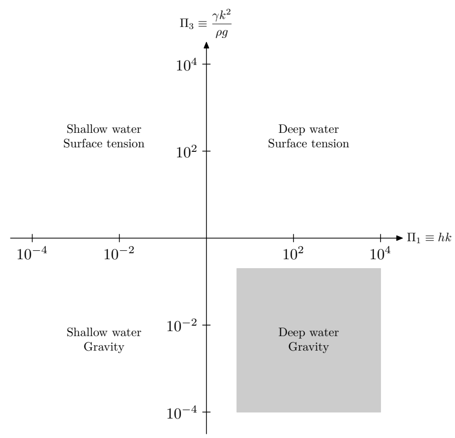

1. Increase g (which is downstairs) by moving to Jupiter.
2. Reduce γ (which is upstairs) by dumping soap on the water.
3. Reduce k (which is upstairs) by studying waves with a huge wavelength.

In this limit, the general deep-water dispersion relation simplifies to

$$\omega^2 = f_{\text{deep}}(0)gk = C_1 gk,$$

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

$$\omega^2 = 1 \times gk,$$

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 [20] were investigating criminal coverups during the Nixon administration, they received help from the mysterious 'Deep Throat', whose valuable advice was to 'follow the money.'

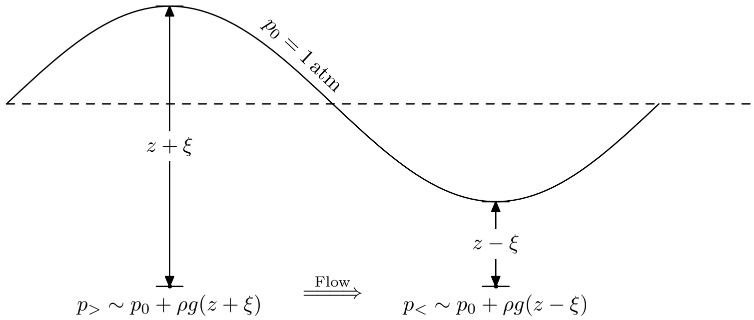

The higher pressure forces water underneath the crest to flow toward the trough, making the water level there rise. Like a swing sliding past equilibrium, the surface overshoots the equilibrium level to produce a new crest and the cycle repeats.

The next step is to quantify the model by estimating sizes, forces, speeds, and energies. In **Section 9.1** we analyzed a messy mortality curve by replacing it with a more tractable shape: a rectangle. The method of discretization 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:

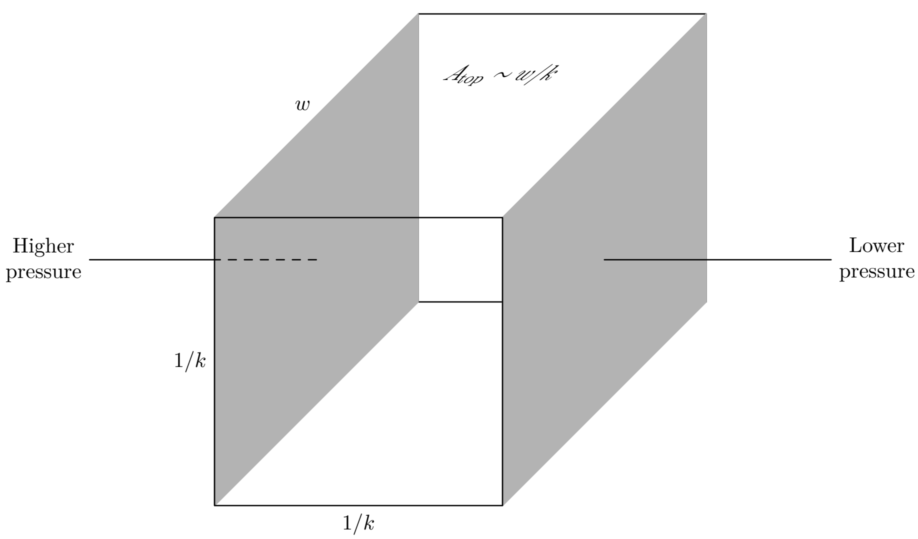

How deep should this slab of water extend? By the Laplace-equation argument, the pressure variation falls off exponentially with depth, with length scale $1/k$. So assume that the slab has a similar length scale, that it has depth $1/k$. What choice do you have? On an infinitely deep ocean, the only 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 ρgh 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

$$F \sim \underbrace{w/k}_{\textit{area}} \times \underbrace{\rho g \xi}_{\Delta p} = \rho g w \xi / k.$$

The slab has mass

$$m = \rho \times \underbrace{w/k^2}_{\textit{volume}},$$

so the force produces an acceleration

$$a_{\text{slab}} \sim \underbrace{\frac{\rho g w \xi}{k}}_{\textit{force}} \bigg/ \underbrace{\frac{\rho w}{k^2}}_{\textit{mass}} = g \xi k.$$

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:

$$a \sim a_{\text{slab}} \sim g \xi k.$$

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 10.1.12**.

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

$$\underbrace{\omega^2\xi}_{a} \sim g\xi k,$$

or

$$\omega^2 = 1\times gk,$$

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.

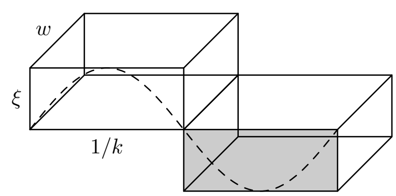

The derivation that resulted in the dispersion relation analyzed the motion of the slab using forces. Another derivation of it uses energy by balancing kinetic and potential energy. To make a wavy surface requires energy, as shown in the figure. The crest rises a characteristic height ξ above the zero of potential, which is the level surface. The volume of water moved upward is $\xi w/k$. So the potential energy is

$$\text{PE}_{\text{gravity}} \sim \underbrace{\rho\xi w/k}_{m}\times g\xi \sim \rho g w\xi^2/k.$$

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

$$\mathrm{KE}_{\mathrm{deep}} \sim \underbrace{\rho w/k^2}_{m_{\mathrm{slab}}} \times \underbrace{\omega^2\xi^2}_{v^2} \sim \rho\omega^2\xi^2 w/k^2.$$

This energy balances the potential energy

$$\underbrace{\rho\omega^2\xi^2 w/k^2}_{KE} \sim \underbrace{\rho g w\xi^2/k}_{PE}.$$

Canceling the factor $\rho w\xi^2$ (in red) common to both energies leaves

$$\omega^2 \sim gk.$$

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:

$$\Pi_2 \sim \frac{\text{kinetic energy in slab}}{\text{gravitational potential energy}} \sim \frac{\omega^2}{gk}.$$

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: $v_{\mathrm{ph}} = w/k$. For deep-water gravity waves, this velocity becomes

$$v_{\mathrm{ph}} = \sqrt{\frac{g}{k}},$$

or, using the dispersion relation to replace k by ω,

$$v_{\mathrm{ph}} = \frac{g}{\omega}.$$

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

$$v_{\text{ph}} = \frac{g}{\omega}.$$

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 10.1.2**. The group velocity is

$$v_{\text{g}} = \frac{\partial}{\partial k}\sqrt{gk} = \frac{1}{2}\sqrt{\frac{g}{k}} = \frac{1}{2}v_{\text{ph}}.$$

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.

10.1.6 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

$$v_{\text{wind}} \sim v_{\text{ph}} \sim \frac{\overbrace{10\,\text{m}\,\text{s}^{-2}}^{g} \times \overbrace{18\,\text{s}}^{T}}{2 \times 3} \sim 30\,\text{m}\,\text{s}^{-1}.$$

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

$$\lambda = v_{\text{ph}}T \sim 30\,\text{m}\,\text{s}^{-1} \times 18\,\text{s} \sim 500\,\text{m}.$$

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 10.1.13**.

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

$$\frac{\text{velocity(18 s waves)}}{\text{velocity(17 s waves)}} = \frac{18}{17} = 1 + \frac{1}{17}.$$

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, $v_g = v_{ph}/2 \sim 15\,\text{m}\,\text{s}^{-1}$, so the storm distance was $d \sim tv_g \sim 2 \cdot 10^4\,\text{km}$, or roughly halfway around the world, an amazingly long and dissipation-free journey.

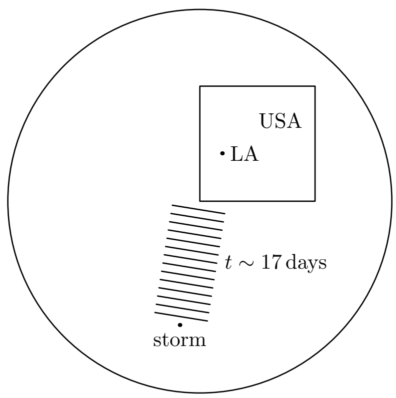

10.1.7 Speedboating

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 $2\,\text{m}\,\text{s}^{-1}$,

$$Re \sim \frac{500\,\text{cm} \times 200\,\text{cm}\,\text{s}^{-1}}{10^{-2}\,\text{cm}^2\,\text{s}^{-1}} \sim 10^7.$$

At such a huge Reynolds number, the flow is turbulent and nearly independent of viscosity (**Section 8.3.7**). 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**:

$$\text{Fr} \equiv \frac{v^2}{gl}.$$

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:

$$v \sim \sqrt{gl}.$$

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

$$v_{\text{boat}}^2 \sim \frac{\omega^2}{k^2} = \frac{g}{k} = g\lambda\!\!\!^{-},$$

where $\lambda\!\!\!^{-} \equiv 1/k = \lambda/2\pi$. 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:

$$\text{Fr} = \frac{v_{\text{boat}}^2/g}{l} \sim \frac{\text{wavelength of bow wave}}{\text{length of boat}}.$$

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 $\lambda\!\!\!^{-}$; if the boat's length l is comparable to $\lambda\!\!\!^{-}$, 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

$$v_{\text{boat}} \sim \sqrt{gl}.$$

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 10.1.7** it determines maximum boating speed. The Froude number is a ratio of potential energy to kinetic energy, as massaging the Froude number shows:

$$\text{Fr} = \frac{v^2}{gl} = \frac{mv^2}{mgl} \sim \frac{\text{kinetic energy}}{\text{potential energy}}.$$

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.

10.1.8 Walking

In the Froude number for walking speed, l is leg length, and gl is a potential energy. For a human with leg length $l \sim 1\,\mathrm{m}$, the condition $\mathrm{Fr} \sim 1$ implies that $v \sim 3\,\mathrm{m\,s^{-1}}$ 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 $4.31\,\mathrm{m\,s^{-1}}$.

This example concludes the study of gravity waves on deep water, which is one corner of the world of waves.

10.1.9 Ripples on deep water

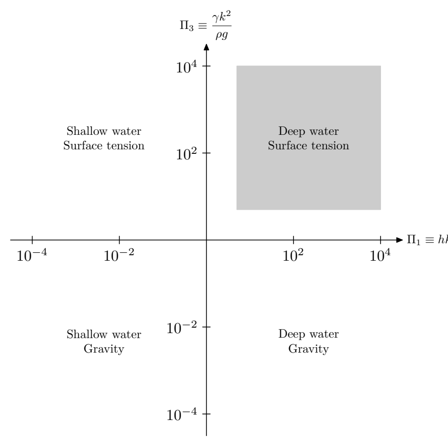

For small wavelengths (large k), surface tension rather than gravity provides the restoring force. This choice brings us to the shaded corner of the figure. If surface tension rather than gravity provides the restoring force, then g vanishes from the final dispersion relation. How to get rid of g and find the new dispersion relation? You could follow the same pattern as for gravity waves (**Section 10.1.5**). 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

$$\omega^2 = 1\times\frac{\gamma k^3}{\rho}.$$

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:

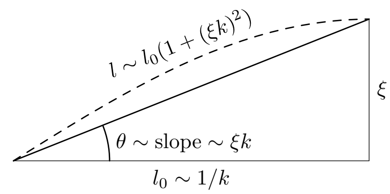

Relative to the level line, the tilted line is longer by a factor $1 + (\xi k)^2$.

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

$$\Delta A \sim \underbrace{w/k}_{\text{level area}} \times \underbrace{(\xi k)^2}_{\text{fractional increase}} \sim w\xi^2 k.$$

The potential energy stored in this extra surface is

$$\text{PE}_{\text{ripple}} \sim \gamma \Delta A \sim \gamma w \xi^2 k.$$

The kinetic energy in the slab is the same as it is for gravity waves, which is:

$$\text{KE} \sim \rho\omega^2\xi^2 w/k^2.$$

Balancing the energies

$$\underbrace{\rho\omega^2\xi^2 w/k^2}_{KE} \sim \underbrace{\gamma w \xi^2 k}_{PE},$$

gives

$$\omega^2 \sim \gamma k^3/\rho.$$

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:

$$\begin{aligned}\Pi_3 &\sim \frac{\text{potential energy in a ripple}}{\text{potential energy in a gravity wave}}\\ &\sim \frac{\gamma w \xi^2 k}{\rho g w \xi^2/k}\\ &\sim \frac{\gamma k^2}{\rho g}.\end{aligned}$$

Alternatively, Π_3 compares $\gamma k^2/\rho$ with g:

$$\Pi_3 \equiv \frac{\gamma k^2/\rho}{g}.$$

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

$$g_{\text{surface tension}} = \gamma k^2/\rho,$$

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 8.3.4**), 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:

$$\underbrace{\rho g w \xi^2/k}_{\textit{PE(gravity wave)}} \xrightarrow{g\to\gamma k^2/\rho} \underbrace{\gamma w \xi^2 k}_{\textit{PE(ripple)}}.$$

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:

$$\omega^2 = gk \xrightarrow{g\to\gamma k^2/\rho} \omega^2 = \frac{\gamma k^3}{\rho}.$$

The interpretation of surface tension as effective gravity is useful when we combine our solutions for gravity waves and for ripples, in **Section 10.1.11** and **Section 10.1.16**. 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

$$v_{\text{ph}} \equiv \frac{\omega}{k} = \sqrt{\frac{\gamma k}{\rho}},$$

and the group velocity is

$$v_{\text{g}} \equiv \frac{\partial\omega}{\partial k} = \frac{3}{2} v_{\text{ph}}.$$

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.

10.1.10 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

$$v_{\mathrm{ph}} = \left(\frac{\overbrace{0.07\,\mathrm{J\,m^{-2}}}^{\gamma} \times \overbrace{600\,\mathrm{m^{-1}}}^{k}}{\underbrace{10^3\,\mathrm{kg\,m^{-3}}}_{\rho}} \right)^{1/2} \sim 21\,\mathrm{cm\,s^{-1}}.$$

According to relation between phase and group velocities, the group velocity is 50 percent larger than the phase velocity: $v_{\mathrm{g}} \sim 30\,\mathrm{cm\,s^{-1}}$. This wavelength of 1 cm is roughly the longest wavelength that still qualifies as a ripple, as shown in an earlier figure repeated here:

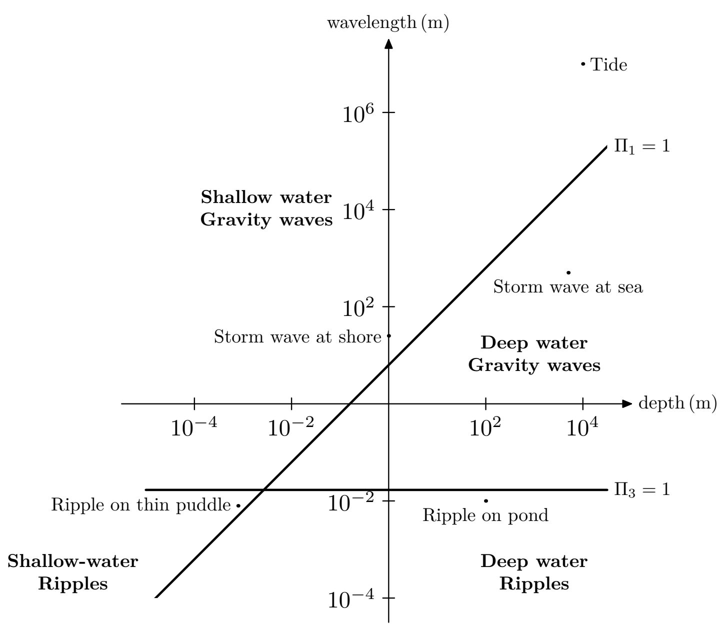

The third dimensionless group, which distinguishes ripples from gravity waves, has value

$$\Pi_3 \equiv \frac{\gamma k^2}{\rho g} \sim \frac{\overbrace{0.07\,\mathrm{J\,m^{-2}}}^{\gamma} \times \overbrace{3.6 \cdot 10^5\,\mathrm{m^{-2}}}^{k^2}}{\underbrace{10^3\,\mathrm{kg\,m^{-3}}}_{\rho} \times \underbrace{10\,\mathrm{m\,s^{-2}}}_{g}} \sim 2.6.$$

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 cm s^{-1} is roughly the minimum phase velocity for ripples. This minimum speed explains why we see mostly $\lambda \sim 1$ 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.

10.1.11 Combining ripples and gravity waves on deep water

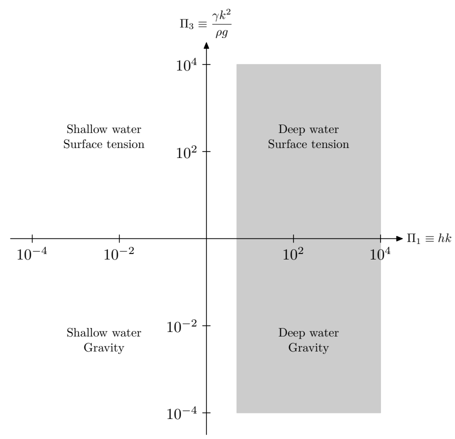

With two corners assembled – gravity waves and ripples in deep water – you can connect the corners to form the deep-water edge. The dispersion relations, for convenience restated here, are

$$\omega^2 = \begin{cases} gk, & \text{gravity waves;} \\ \gamma k^3/\rho, & \text{ripples.} \end{cases}$$

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:

$$\omega^2 = gk + \gamma k^3/\rho.$$

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

$$g_{\text{total}} = g + g_{\text{surface tension}} = g + \frac{\gamma k^2}{\rho}.$$

Replace g by g_{total} in $\omega^2 = gk$ reproduces the deep-water dispersion relation:

$$\omega^2 = \left(g + \frac{\gamma k^2}{\rho}\right)k = gk + \gamma k^3/\rho.$$

This dispersion relation tells us wave speeds for all wavelengths or wavenumbers. The phase velocity is

$$v_{\text{ph}} \equiv \frac{\omega}{k} = \sqrt{\frac{\gamma k}{\rho} + \frac{g}{k}}.$$

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:

$$\frac{\gamma k_0}{\rho} \sim \frac{g}{k_0}.$$

This computation is an example of order-of-magnitude minimization. The minimum-speed wavenumber is

$$k_0 \sim \sqrt{\frac{\rho g}{\gamma}}.$$

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

$$v_0 \sim \sqrt{\frac{2g}{k_0}} \sim \left(\frac{4\gamma g}{\rho}\right)^{1/4}.$$

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

$$v_0 \sim 23\,\mathrm{cm\,s}^{-1}.$$

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

$$v_{\mathrm{g}} \approx 17.7\,\mathrm{cm\,s}^{-1}.$$

[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

$$\rho = \gamma = g = 1.$$

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

$$v_{\text{ph}}^2 = k + \frac{1}{k}.$$

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,

$$\underbrace{(a+b)/2}_{AM} \geq \underbrace{\sqrt{ab}}_{GM},$$

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

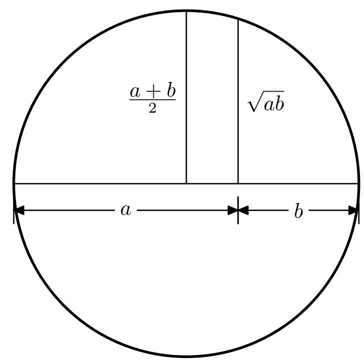

$$k + \frac{1}{k} \geq 2,$$

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

$$\omega = \sqrt{k^3 + k},$$

and the group velocity is

$$v_{\text{g}} = \frac{\partial}{\partial k}\sqrt{k^3 + k},$$

which is

$$v_{\text{g}} = \frac{1}{2}\frac{3k^2 + 1}{\sqrt{k^3 + k}}.$$

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

$$\frac{\partial}{\partial k}\underbrace{\frac{9k^4+6k^2+1}{k^3+k}}_{v_g^2}=\frac{(3k^2+1)(3k^4+6k^2-1)}{(k^3+k)^2}.$$

Equating this derivative to zero gives $3k^4+6k^2-1=0$, which is a quadratic in k^2, and has positive solution

$$k_1=\sqrt{-1+\sqrt{4/3}}\sim 0.393.$$

At this k, the group velocity is

$$v_g(k_1)\approx 1.086.$$

In more usual units, this minimum velocity is

$$v_g\approx 1.086\left(\frac{\gamma g}{\rho}\right)^{1/4}.$$

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.

10.1.12 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.

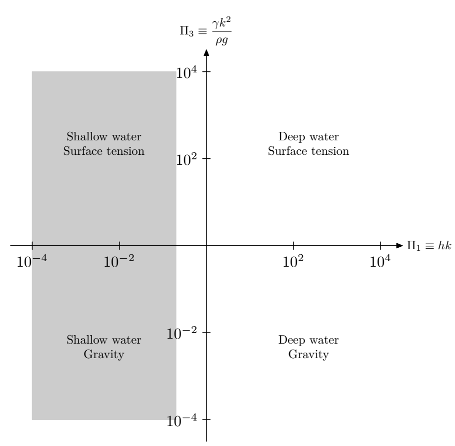

In deep water the slab has depth $1/k$. In shallow water, however, where $h \ll 1/k$, the bottom of the ocean arrives before that depth. So the shallow-water slab has depth h. Its length is still $1/k$, and its width is still w. Because the depth changed, the argument about how the water flows is slightly 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

$$\mathrm{KE}_{\text{shallow}} \sim \underbrace{\rho wh/k}_{\textit{mass}} \times \underbrace{(\xi\omega/hk)^2}_{v^2} \sim \frac{\rho w \xi^2 \omega^2}{hk^3}.$$

This energy balances the potential energy, a computation we do for the two limiting cases: ripples and gravity waves.

10.1.13 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

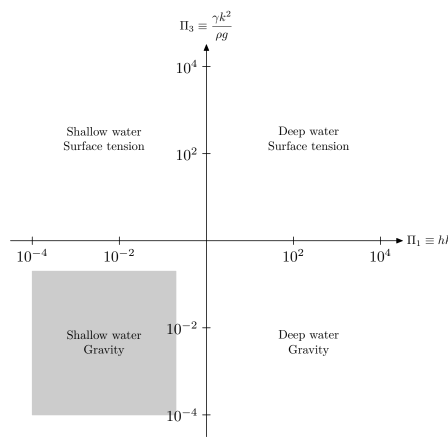

$$\mathrm{PE} \sim \rho g w \xi^2/k.$$

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 10.1.15**, the water depth determines the kinetic energy.]

Balancing this energy against the kinetic energy gives:

$$\underbrace{\frac{\rho w \xi^2 \omega^2}{hk^3}}_{KE} \sim \underbrace{\rho g w \xi^2/k}_{PE}.$$

So

$$\omega^2 = 1 \times ghk^2.$$

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

$$f_{\text{shallow}}(kh, \frac{\gamma k^2}{\rho g}) = kh.$$

Since $\omega \propto k^1$, the group and phase velocities are equal and independent of frequency:

$$v_{\text{ph}} = \frac{\omega}{k} = \sqrt{gh},$$
$$v_{\text{g}} = \frac{\partial \omega}{\partial k} = \sqrt{gh}.$$

Shallow water is **nondispersive**: All frequencies move at the same velocity, so pulses composed of various frequencies propagate without smearing.

10.1.14 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.

10.1.15 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

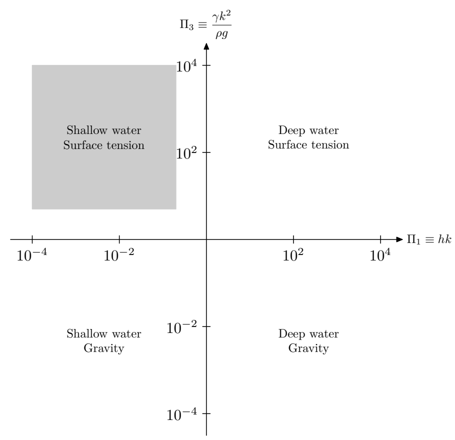

$$\mathrm{PE}_{\mathrm{ripple}} \sim \gamma \Delta A \sim \gamma w \xi^2 k.$$

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:

$$\underbrace{\frac{\rho w \xi^2 \omega^2}{hk^3}}_{KE} \sim \underbrace{\frac{w}{k}\gamma(k\xi)^2}_{PE}.$$

Then

$$\omega^2 \sim \frac{\gamma h k^4}{\rho}.$$

The phase velocity is

$$v_{\mathrm{ph}} = \frac{\omega}{k} = \sqrt{\frac{\gamma h k^2}{\rho}},$$

and the group velocity is $v_{\mathrm{g}} = 2v_{\mathrm{ph}}$ (the form of the dispersion relation is $\omega \propto k^2$). For $h \sim 1$ mm, this speed is

$$v \sim \left(\frac{0.07\,\mathrm{N\,m^{-1}} \times 10^{-3}\,\mathrm{m} \times 3.6 \cdot 10^{5}\,\mathrm{m^{-2}}}{10\,\mathrm{kg\,m^{-3}}} \right)^{1/2} \sim 16\,\mathrm{cm\,s^{-1}}.$$

10.1.16 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

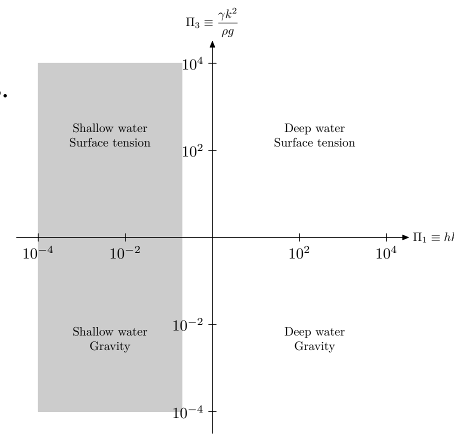

$$\omega^2 \sim k^2 \left(gh + \frac{\gamma h k^2}{\rho} \right).$$

10.1.17 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:

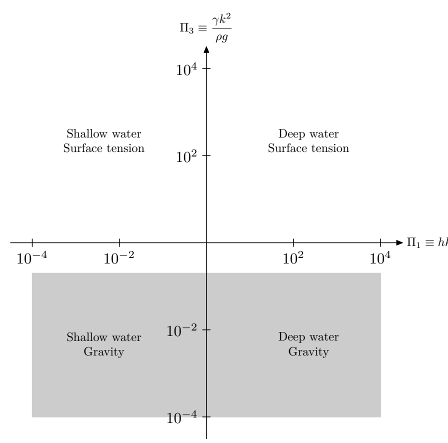

$$\omega^2 = gk \times \begin{cases} 1, & \text{deep water;} \\ hk, & \text{shallow water.} \end{cases}$$

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:

$$\omega^2 = gk \tanh hk.$$

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.

10.1.18 Combining deep- and shallow-water ripples

We now examine the final edge: ripples in shallow and deep water, as shown in the figure. In **Section 10.1.17**, $\tanh kh$ 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:

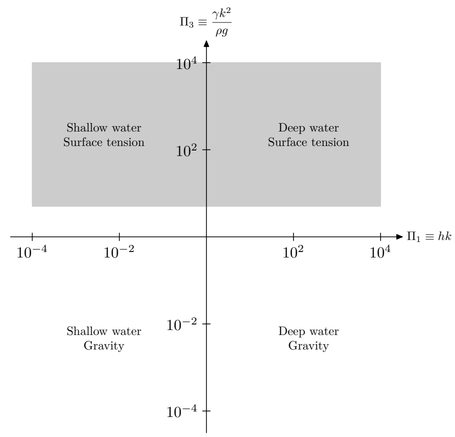

$$\omega^2 = \begin{cases} \gamma k^3/\rho, & \text{if } kh \gg 1; \\ \gamma hk^4/\rho, & \text{if } kh \ll 1. \end{cases}$$

If we factor out $\gamma k^3/\rho$, the necessary transformation becomes clear:

$$\omega^2 = \frac{\gamma k^3}{\rho} \times \begin{cases} 1, & \text{if } kh \gg 1; \\ hk, & \text{if } kh \ll 1. \end{cases}$$

This ripple result looks similar to the gravity-wave result, so make the same replacement:

$$\begin{cases} 1, & \text{if } kh \gg 1, \\ hk, & \text{if } kh \ll 1, \end{cases} \quad \text{becomes} \quad \tanh kh.$$

Then you get the general ripple dispersion relation:

$$\omega^2 = \frac{\gamma k^3}{\rho} \tanh kh.$$

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.

10.1.19 Combining all the analyses

Now we can replace g with g_{total}, to find the One True Dispersion Relation:

$$\omega^2 = (gk + \gamma k^3/\rho) \tanh kh.$$

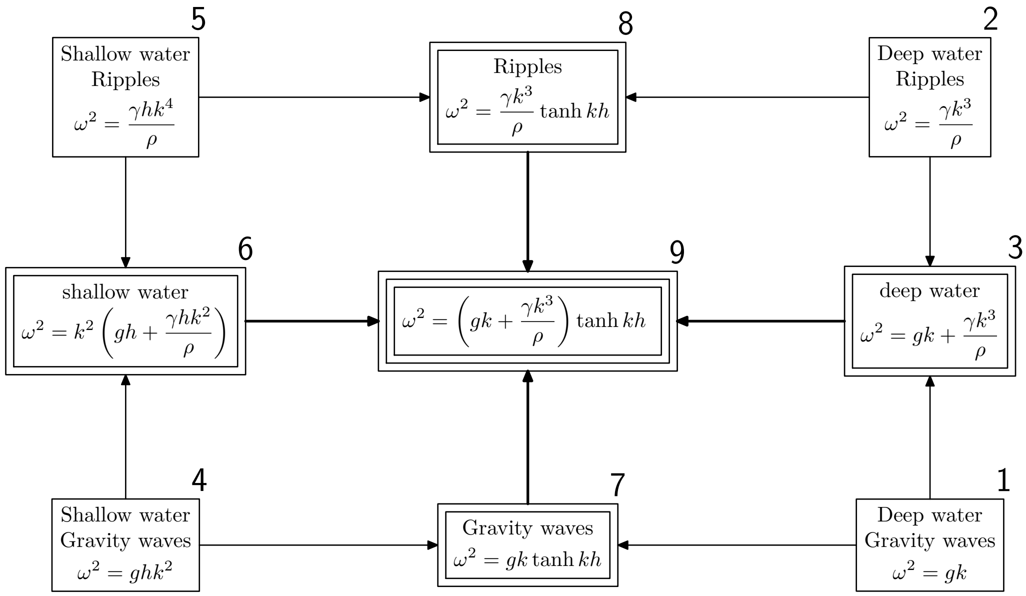

Each box in the figure represents a special case. The numbers next to the boxes mark the order in which we studied that limit. In the final step (9), we combined all the analyses into the superbox in the center, which contains the dispersion relation for all waves: gravity waves or ripples, shallow water 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.

10.1.20 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 discretization 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.

Exercises

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?

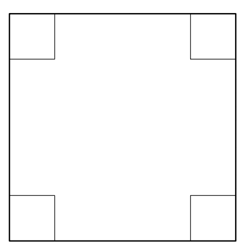

Minima without calculus.

Impossible

How can tidal waves on the ocean (typical depth $\sim 4\,\text{km}$) be considered shallow water?

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?

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?

Kelvin wedge

Show that the opening angle in a ship wake is $2\sin^{-1}(1/3)$.

Part 4

Backmatter

Chapter 11
Bon voyage!

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.

11.1 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.

11.2 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 8.1**.

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:

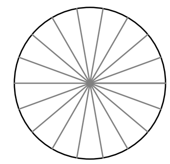

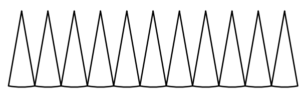

Each pie wedge is almost a triangle, so its area is $bh/2$, where the height h is approximately r. The sum of all the bases is the circumference $2\pi r$, so $A = 2\pi r \times r/2 = \pi r^2$.

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.

11.3 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?

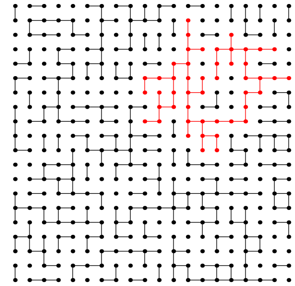

cluster = 12%
$p = 0.40$

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.

Here is what happens as p_{bond} increases from 0.40 to 0.50 to 0.55 to 0.60:

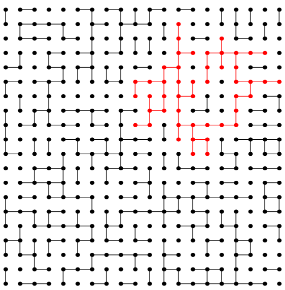

cluster = 12%
$p = 0.40$

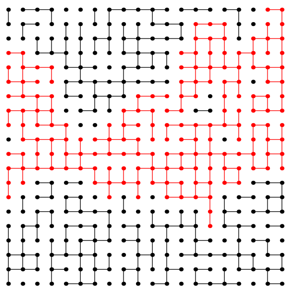

cluster = 43%
$p = 0.50$

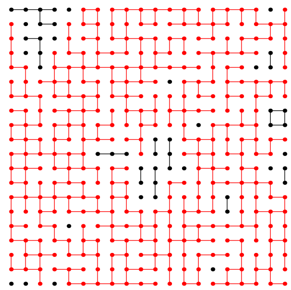

cluster = 87%
$p = 0.55$

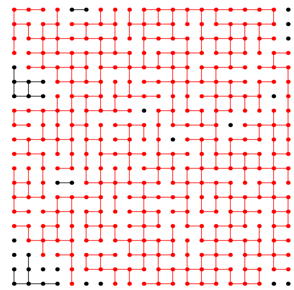

cluster = 91%
$p = 0.60$

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.

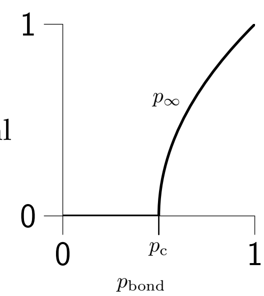

An analogy to learning is that each lattice point (each dot) is a fact or formula, and each bond links two facts. For long-lasting learning, you want the facts to support each other via their connections. Let's say that you want the facts to become part of an infinite and therefore self-supporting lattice. 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] Mike Gancarz. *Linux and the Unix Philosophy*. Digital Press, 2nd revised edition, 2002.

[2] Brian W. Kernighan and Rob Pike. *The Unix Programming Environment*. Prentice Hall, 1984.

[3] Eric S. Raymond. *The Art of UNIX Programming*. Addison-Wesley, 2003.

[4] Knut Schmid-Nielsen. *Scaling: Why Animal Size is So Important*. Cambridge University Press, Cambridge, England, 1984.

[5] R. David Middlebrook. Methods of design-oriented analysis: Low-entropy expressions. In *New Approaches to Undergraduate Engineering Education IV*, 1992. See `mit.edu/6.969/www/`.

[6] .

[7] Richard Feynman, Robert B. Leighton and Matthew Sands. *The Feynman Lectures on Physics*, volume I. Addison-Wesley, Reading, MA, 1963a.

[8] 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.

[9] E. Buckingham. On physically similar systems. *Physical Review*, 4(4):345–376, 1914.

[10] S. Chandrasekhar. *An Introduction to the Study of Stellar Structure*. University of Chicago Press, Chicago, Ill., 1939.

[11] S. Chandrasekhar. *Hydrodynamic and Hydromagnetic Stability*. Oxford University Press, Oxford, 1961.

[12] Horace Lamb. *Hydrodynamics*. Dover, New York, 6th edition, 1932/1993.

[13] Steven Vogel. *Life in Moving Fluids: The Physical Biology of Flow*. Princeton University Press, Princeton, NJ, 1994.

[14] E. M. Purcell. Life at low Reynolds number. *American Journal of Physics*, 45:3–11, 1977.

[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] Albert Einstein. *The principle of relativity*. Dover Publications, New York, 1952 [1923].

[17] Albert Einstein. *Relativity: The special and the general theory*. Three Rivers Press, reprint edition, 1995.

[18] T. E. Faber. *Fluid Dynamics for Physicists*. Cambridge University Press, Cambridge, England, 1995.

[19] Horace Lamb. *Hydrodynamics*. Dover, New York, Sixth edition, 1945.

[20] Carl Bernstein and Bob Woodward. *All the President's Men*. Simon and Schuster, New York, 1974.

Index