


Part III

Technical chapters



A Cars II

We estimated that a car driven 100 km uses about 80 kWh of energy.

Where does this energy go? How does it depend on properties of the car? Could we make cars that are 100 times more efficient? Let's make a simple cartoon of car-driving, to describe where the energy goes. The energy in a typical fossil-fuel car goes to four main destinations, all of which we will explore:

1. speeding up then slowing down using the brakes;
2. air resistance;
3. rolling resistance;
4. heat – 75% of the energy is thrown away as heat, because the energy-conversion chain is inefficient.

Initially our cartoon will ignore rolling resistance; we'll add in this effect later in the chapter.

Assume the driver accelerates rapidly up to a cruising speed v , and maintains that speed for a distance d , which is the distance between traffic lights, stop signs, or congestion events. At this point, he slams on the brakes and turns all his kinetic energy into heat in the brakes. (This vehicle doesn't have fancy regenerative braking.) Once he's able to move again, he accelerates back up to his cruising speed, v . This acceleration gives the car kinetic energy; braking throws that kinetic energy away.

Energy goes not only into the brakes: while the car is moving, it makes air swirl around. A car leaves behind it a tube of swirling air, moving at a speed similar to v . Which of these two forms of energy is the bigger: kinetic energy of the swirling air, or heat in the brakes? Let's work it out.

- The car speeds up and slows down once in each duration d/v . The rate at which energy pours into the brakes is:

$$\frac{\text{kinetic energy}}{\text{time between braking events}} = \frac{\frac{1}{2}m_c v^2}{d/v} = \frac{\frac{1}{2}m_c v^3}{d}, \quad (\text{A.1})$$

where m_c is the mass of the car.




Figure A.1. A Peugeot 206 has a drag coefficient of 0.33. Photo by Christopher Batt.

The key formula for most of the calculations in this book is:

$$\text{kinetic energy} = \frac{1}{2}mv^2.$$

For example, a car of mass $m = 1000$ kg moving at 100 km per hour or $v = 28$ m/s has an energy of

$$\frac{1}{2}mv^2 \simeq 390\,000\text{ J} \simeq 0.1\text{ kWh}.$$




Figure A.2. Our cartoon: a car moves at speed v between stops separated by a distance d .

Figure A.3. A car moving at speed v creates behind it a tube of swirling air; the cross-sectional area of the tube is similar to the frontal area of the car, and the speed at which air in the tube swirls is roughly v .

- The tube of air created in a time t has a volume Avt , where A is the cross-sectional area of the tube, which is similar to the area of the front view of the car. (For a streamlined car, A is usually a little smaller than the frontal area A_{car} , and the ratio of the tube's effective cross-sectional area to the car area is called the drag coefficient c_d . Throughout the following equations, A means the effective area of the car, $c_d A_{\text{car}}$.) The tube has mass $m_{\text{air}} = \rho Avt$ (where ρ is the density of air) and swirls at speed v , so its kinetic energy is:

$$\frac{1}{2} m_{\text{air}} v^2 = \frac{1}{2} \rho Avt v^2,$$

and the rate of generation of kinetic energy in swirling air is:

$$\frac{\frac{1}{2} \rho Avt v^2}{t} = \frac{1}{2} \rho A v^3.$$

So the total rate of energy production by the car is:

$$\begin{aligned} \text{power going into brakes} &+ \text{power going into swirling air} \\ = \frac{1}{2} m_c v^3 / d &+ \frac{1}{2} \rho A v^3. \end{aligned} \quad (\text{A.2})$$

Both forms of energy dissipation scale as v^3 . So this cartoon predicts that a driver who halves his speed v makes his power consumption 8 times smaller. If he ends up driving the same total distance, his journey will take twice as long, but the total energy consumed by his journey will be four times smaller.

Which of the two forms of energy dissipation – brakes or air-swirling – is the bigger? It depends on the ratio of

$$(m_c / d) / (\rho A).$$

If this ratio is much bigger than 1, then more power is going into brakes; if it is smaller, more power is going into swirling air. Rearranging this ratio, it is bigger than 1 if

$$m_c > \rho A d.$$

Now, Ad is the volume of the tube of air swept out from one stop sign to the next. And ρAd is the mass of that tube of air. So we have a very simple situation: energy dissipation is dominated by kinetic-energy-being-dumped-into-the-brakes if the mass of the car is *bigger* than the mass of the tube of air from one stop sign to the next; and energy dissipation is dominated by making-air-swirl if the mass of the car is *smaller* (figure A.4).

Let's work out the special distance d^* between stop signs, below which the dissipation is braking-dominated and above which it is air-swirling dominated (also known as drag-dominated). If the frontal area of the car is:

$$A_{\text{car}} = 2 \text{ m wide} \times 1.5 \text{ m high} = 3 \text{ m}^2$$

I'm using this formula:

$$\text{mass} = \text{density} \times \text{volume}$$

The symbol ρ (Greek letter 'rho') denotes the density.




Figure A.4. To know whether energy consumption is braking-dominated or air-swirling-dominated, we compare the mass of the car with the mass of the tube of air between stop-signs.



Figure A.5. Power consumed by a car is proportional to its cross-sectional area, during motorway driving, and to its mass, during town driving. Guess which gets better mileage – the VW on the left, or the spaceship?

and the drag coefficient is $c_d = 1/3$ and the mass is $m_c = 1000\text{ kg}$ then the special distance is:

$$d^* = \frac{m_c}{\rho c_d A_{\text{car}}} = \frac{1000\text{ kg}}{1.3\text{ kg/m}^3 \times \frac{1}{3} \times 3\text{ m}^2} = 750\text{ m}.$$

So “city-driving” is dominated by kinetic energy and braking if the distance between stops is less than 750 m. Under these conditions, it’s a good idea, if you want to save energy:

- 1. to reduce the mass of your car;
- 2. to get a car with regenerative brakes (which roughly halve the energy lost in braking – see Chapter 20); and
- 3. to drive more slowly.

When the stops are significantly more than 750m apart, energy dissipation is drag-dominated. Under these conditions, it doesn’t much matter what your car weighs. Energy dissipation will be much the same whether the car contains one person or six. Energy dissipation can be reduced:

- 1. by reducing the car’s drag coefficient;
- 2. by reducing its cross-sectional area; or
- 3. by driving more slowly.

The actual energy consumption of the car will be the energy dissipation in equation (A.2), cranked up by a factor related to the inefficiency of the engine and the transmission. Typical petrol engines are about 25% efficient, so of the chemical energy that a car guzzles, three quarters is wasted in making the car’s engine and radiator hot, and just one quarter goes into “useful” energy:

$$\text{total power of car} \simeq 4 \left[\frac{1}{2} m_c v^3 / d + \frac{1}{2} \rho A v^3 \right].$$

Let’s check this theory of cars by plugging in plausible numbers for motorway driving. Let $v = 70\text{ miles per hour} = 110\text{ km/h} = 31\text{ m/s}$ and $A = c_d A_{\text{car}} = 1\text{ m}^2$. The power consumed by the engine is estimated to be roughly

$$4 \times \frac{1}{2} \rho A v^3 = 2 \times 1.3\text{ kg/m}^3 \times 1\text{ m}^2 \times (31\text{ m/s})^3 = 80\text{ kW}.$$

If you drive the car at this speed for one hour every day, then you travel 110km and use **80 kWh** of energy per day. If you drove at half this speed for two hours per day instead, you would travel the same distance and use up **20 kWh** of energy. This simple theory seems consistent with the

ENERGY-PER-DISTANCE	
Car at 110 km/h	↔ 80 kWh/(100 km)
Bicycle at 21 km/h	↔ 2.4 kWh/(100 km)

PLANES AT 900 KM/H	
A380	27 kWh/100 seat-km

Table A.6. Facts worth remembering: car energy consumption.

mileage figures for cars quoted in Chapter 3. Moreover, the theory gives insight into how the energy consumed by your car could be reduced. The theory has a couple of flaws which we'll explore in a moment.

Could we make a new car that consumes 100 times less energy and still goes at 70 mph? **No**. Not if the car has the same shape. On the motorway at 70 mph, the energy is going mainly into making air swirl. Changing the materials the car is made from makes no difference to that. A miraculous improvement to the fossil-fuel engine could perhaps boost its efficiency from 25% to 50%, bringing the energy consumption of a fossil-fuelled car down to roughly 40 kWh per 100 km.

Electric vehicles have some wins: while the weight of the energy store, per useful kWh stored, is about 25 times bigger than that of petrol, the weight of an electric engine can be about 8 times smaller. And the energy-chain in an electric car is much more efficient: electric motors can be 90% efficient.

We'll come back to electric cars in more detail towards the end of this chapter.

Bicycles and the scaling trick

Here's a fun question: what's the energy consumption of a bicycle, in kWh per 100 km? Pushing yourself along on a bicycle requires energy for the same reason as a car: you're making air swirl around. Now, we could do all the calculations from scratch, replacing car-numbers by bike-numbers. But there's a simple trick we can use to get the answer for the bike from the answer for the car. The energy consumed by a car, per distance travelled, is the power-consumption associated with air-swirling,

$$4 \times \frac{1}{2} \rho A v^3,$$

divided by the speed, v ; that is,

$$\text{energy per distance} = 4 \times \frac{1}{2} \rho A v^2.$$

The "4" came from engine inefficiency; ρ is the density of air; the area $A = c_d A_{\text{car}}$ is the effective frontal area of a car; and v is its speed.

Now, we can compare a bicycle with a car by dividing $4 \times \frac{1}{2} \rho A v^2$ for the bicycle by $4 \times \frac{1}{2} \rho A v^2$ for the car. All the fractions and ρ s cancel, if the efficiency of the carbon-powered bicyclist's engine is similar to the efficiency of the carbon-powered car engine (which it is). The ratio is:

$$\frac{\text{energy per distance of bike}}{\text{energy per distance of car}} = \frac{c_d^{\text{bike}} A_{\text{bike}} v_{\text{bike}}^2}{c_d^{\text{car}} A_{\text{car}} v_{\text{car}}^2}.$$

The trick we are using is called "scaling." If we know how energy consumption scales with speed and area, then we can predict energy con-

DRAG COEFFICIENTS	
CARS	
Honda Insight	0.25
Prius	0.26
Renault 25	0.28
Honda Civic (2006)	0.31
VW Polo GTi	0.32
Peugeot 206	0.33
Ford Sierra	0.34
Audi TT	0.35
Honda Civic (2001)	0.36
Citroën 2CV	0.51
Cyclist	0.9
Long-distance coach	0.425
PLANES	
Cessna	0.027
Learjet	0.022
Boeing 747	0.031
DRAG-AREAS (m ²)	
Land Rover Discovery	1.6
Volvo 740	0.81
Typical car	0.8
Honda Civic	0.68
VW Polo GTi	0.65
Honda Insight	0.47

Table A.7. Drag coefficients and drag areas.

sumption of objects with completely different speeds and areas. Specifically, let’s assume that the area ratio is

$$\frac{A_{\text{bike}}}{A_{\text{car}}} = \frac{1}{4}.$$

(Four cyclists can sit shoulder to shoulder in the width of one car.) Let’s assume the bike is not very well streamlined:

$$\frac{c_{\text{d}}^{\text{bike}}}{c_{\text{d}}^{\text{car}}} = \frac{1}{3}$$

And let’s assume the speed of the bike is 21 km/h (13 miles per hour), so

$$\frac{v_{\text{bike}}}{v_{\text{car}}} = \frac{1}{5}.$$

Then

$$\begin{aligned} \frac{\text{energy-per-distance of bike}}{\text{energy-per-distance of car}} &= \left(\frac{c_{\text{d}}^{\text{bike}}}{c_{\text{d}}^{\text{car}}} \frac{A_{\text{bike}}}{A_{\text{car}}} \right) \left(\frac{v_{\text{bike}}}{v_{\text{car}}} \right)^2 \\ &= \left(\frac{3}{4} \right) \times \left(\frac{1}{5} \right)^2 \\ &= \frac{3}{100} \end{aligned}$$

So a cyclist at 21 km/h consumes about 3% of the energy per kilometre of a lone car-driver on the motorway – about **2.4 kWh per 100 km**.

If you would like a vehicle whose fuel efficiency is 30 times better than a car’s, it’s simple: ride a bike.

What about rolling resistance?

Some things we’ve completely ignored so far are the energy consumed in the tyres and bearings of the car, the energy that goes into the noise of wheels against asphalt, the energy that goes into grinding rubber off the tyres, and the energy that vehicles put into shaking the ground. Collectively, these forms of energy consumption are called *rolling resistance*. The standard model of rolling resistance asserts that the force of rolling resistance is simply proportional to the weight of the vehicle, independent of

wheel	C_{rr}
train (steel on steel)	0.002
bicycle tyre	0.005
truck rubber tyres	0.007
car rubber tyres	0.010

Table A.8. The rolling resistance is equal to the weight multiplied by the coefficient of rolling resistance, C_{rr} . The rolling resistance includes the force due to wheel flex, friction losses in the wheel bearings, shaking and vibration of both the roadbed and the vehicle (including energy absorbed by the vehicle’s shock absorbers), and sliding of the wheels on the road or rail. The coefficient varies with the quality of the road, with the material the wheel is made from, and with temperature. The numbers given here assume smooth roads. [2bhu35]




Figure A.9. Simple theory of car fuel consumption (energy per distance) when driving at steady speed. Assumptions: the car's engine uses energy with an efficiency of 0.25, whatever the speed; $c_d A_{\text{car}} = 1 \text{ m}^2$; $m_{\text{car}} = 1000 \text{ kg}$; and $C_{\text{rr}} = 0.01$.




Figure A.10. Simple theory of bike fuel consumption (energy per distance). Vertical axis is energy consumption in kWh per 100 km. Assumptions: the bike's engine (that's you!) uses energy with an efficiency of 0.25; the drag-area of the cyclist is 0.75 m^2 ; the cyclist+bike's mass is 90 kg ; and $C_{\text{rr}} = 0.005$.




Figure A.11. Simple theory of train energy consumption, *per passenger*, for an eight-carriage train carrying 584 passengers. Vertical axis is energy consumption in kWh per 100 p-km. Assumptions: the train's engine uses energy with an efficiency of 0.90; $c_d A_{\text{train}} = 11 \text{ m}^2$; $m_{\text{train}} = 400\,000 \text{ kg}$; and $C_{\text{rr}} = 0.002$.

the speed. The constant of proportionality is called the coefficient of rolling resistance, C_{rr} . Table A.8 gives some typical values.

The coefficient of rolling resistance for a car is about 0.01. The effect of rolling resistance is just like perpetually driving up a hill with a slope of one in a hundred. So rolling friction is about 100 newtons per ton, independent of speed. You can confirm this by pushing a typical one-ton car along a flat road. Once you've got it moving, you'll find you can keep it moving with one hand. (100 newtons is the weight of 100 apples.) So at a speed of 31 m/s (70 mph), the power required to overcome rolling resistance, for a one-ton vehicle, is

$$\text{force} \times \text{velocity} = (100 \text{ newtons}) \times (31 \text{ m/s}) = 3100 \text{ W};$$

which, allowing for an engine efficiency of 25%, requires 12 kW of power to go into the engine; whereas the power required to overcome drag was estimated on p256 to be 80 kW. So, at high speed, about 15% of the power is required for rolling resistance.

Figure A.9 shows the theory of fuel consumption (energy per unit distance) as a function of steady speed, when we add together the air resistance and rolling resistance.

The speed at which a car's rolling resistance is equal to air resistance is

given by

$$C_{rr}m_cg = \frac{1}{2}\rho c_d A v^2,$$

that is,

$$v = \sqrt{2 \frac{C_{rr}m_cg}{\rho c_d A}} = 7 \text{ m/s} = 16 \text{ miles per hour.}$$

Bicycles

For a bicycle ($m = 90 \text{ kg}$, $A = 0.75 \text{ m}^2$), the transition from rolling-resistance-dominated cycling to air-resistance-dominated cycling takes place at a speed of about 12 km/h. At a steady speed of 20 km/h, cycling costs about **2.2 kWh per 100 km**. By adopting an aerodynamic posture, you can reduce your drag area and cut the energy consumption down to about 1.6 kWh per 100 km.

Trains

For an eight-carriage train as depicted in figure 20.4 ($m = 400\,000 \text{ kg}$, $A = 11 \text{ m}^2$), the speed above which air resistance is greater than rolling resistance is

$$v = 33 \text{ m/s} = 74 \text{ miles per hour.}$$

For a single-carriage train ($m = 50\,000 \text{ kg}$, $A = 11 \text{ m}^2$), the speed above which air resistance is greater than rolling resistance is

$$v = 12 \text{ m/s} = 26 \text{ miles per hour.}$$

Dependence of power on speed

When I say that halving your driving speed should reduce fuel consumption (in miles per gallon) to *one quarter* of current levels, some people feel sceptical. They have a point: most cars' engines have an optimum revolution rate, and the choice of gears of the car determines a range of speeds at which the optimum engine efficiency can be delivered. If my suggested experiment of halving the car's speed takes the car out of this designed range of speeds, the consumption might not fall by as much as four-fold. My tacit assumption that the engine's efficiency is the same at all speeds and all loads led to the conclusion that it's always good (in terms of miles per gallon) to travel slower; but if the engine's efficiency drops off at low speeds, then the most fuel-efficient speed might be at an intermediate speed that makes a compromise between going slow and keeping the engine efficient. For the BMW 318ti in figure A.12, for example, the optimum speed is about 60 km/h. But if society were to decide that car speeds should be reduced, there is nothing to stop engines and gears being redesigned so that the peak engine efficiency was found at the right speed. As further evidence




Figure A.12. Current cars' fuel consumptions do not vary as speed squared. Prius data from B.Z. Wilson; BMW data from Phil C. Stuart. The smooth curve shows what a speed-squared curve would look like, assuming a drag-area of 0.6 m^2 .




Figure A.13. Powers of cars (kW) versus their top speeds (km/h). Both scales are logarithmic. The power increases as the third power of the speed. To go twice as fast requires eight times as much engine power. From Tennekes (1997).

that the power a car requires really does increase as the cube of speed, figure A.13 shows the engine power versus the top speeds of a range of cars. The line shows the relationship “power proportional to v^3 .”

Electric cars: is range a problem?

People often say that the range of electric cars is not big enough. Electric car advocates say “no problem, we can just put in more batteries” – and that’s true, but we need to work out what effect the extra batteries have on the energy consumption. The answer depends sensitively on what energy density we assume the batteries deliver: for an energy density of 40 Wh/kg (typical of lead-acid batteries), we’ll see that it’s hard to push the range beyond 200 or 300 km; but for an energy density of 120 Wh/kg (typical of various lithium-based batteries), a range of 500 km is easily achievable.

Let’s assume that the mass of the car and occupants is 740 kg, *without* any batteries. In due course we’ll add 100 kg, 200 kg, 500 kg, or perhaps 1000 kg of batteries. Let’s assume a typical speed of 50 km/h (30 mph); a drag-area of 0.8 m^2 ; a rolling resistance of 0.01; a distance between stops of 500 m; an engine efficiency of 85%; and that during stops and starts, regenerative braking recovers half of the kinetic energy of the car. Charging up the car from the mains is assumed to be 85% efficient. Figure A.14 shows the transport cost of the car versus its range, as we vary the amount of battery on board. The upper curve shows the result for a battery whose energy density is 40 Wh/kg (old-style lead-acid batteries). The range is limited by a wall at about 500 km. To get close to this maximum range, we have to take along comically large batteries: for a range of 400 km, for example, 2000 kg of batteries are required, and the transport cost is above 25 kWh per 100 km. If we are content with a range of 180 km, however, we can get by with 500 kg of batteries. Things get much better when we switch to lighter lithium-ion batteries. At an energy density of 120 Wh/kg, electric cars with 500 kg of batteries can easily deliver a range of 500 km. The transport cost is predicted to be about 13 kWh per 100 km.

It thus seems to me that the range problem has been solved by the advent of modern batteries. It would be nice to have even better batteries, but an energy density of 120 Wh per kg is already good enough, as long as we’re happy for the batteries in a car to weigh up to 500 kg. In practice I imagine most people would be content to have a range of 300 km, which can be delivered by 250 kg of batteries. If these batteries were divided into ten 25 kg chunks, separately unpluggable, then a car user could keep just four of the ten chunks on board when he’s doing regular commuting (100 kg gives a range of 140 km); and collect an extra six chunks from a battery-recharging station when he wants to make longer-range trips. During long-range trips, he would exchange his batteries for a fresh set at a battery-exchange station every 300 km or so.




Figure A.14. Theory of electric car range (horizontal axis) and transport cost (vertical axis) as a function of battery mass, for two battery technologies. A car with 500 kg of old batteries, with an energy density of 40 Wh per kg, has a range of 180 km. With the same weight of modern batteries, delivering 120 Wh per kg, an electric car can have a range of more than 500 km. Both cars would have an energy cost of about 13 kWh per 100 km. These numbers allow for a battery charging efficiency of 85%.

Notes and further reading

page no.

256 *Typical petrol engines are about 25% efficient.* Encarta [6by8x] says “The efficiencies of good modern Otto-cycle engines range between 20 and 25%.” The petrol engine of a Toyota Prius, famously one of the most efficient car engines, uses the Atkinson cycle instead of the Otto cycle; it has a peak power output of 52 kW and has an efficiency of 34% when delivering 10 kW [348whs]. The most efficient diesel engine in the world is 52%-efficient, but it’s not suitable for cars as it weighs 2300 tons: the Wartsila–Sulzer RTA96-C turbocharged diesel engine (figure A.15) is intended for container ships and has a power output of 80 MW.

– *Regenerative brakes roughly halve the energy lost in braking.* Source: E4tech (2007).

257 *Electric engines can be about 8 times lighter than petrol engines.*

A 4-stroke petrol engine has a power-to-mass ratio of roughly 0.75 kW/kg. The best electric motors have an efficiency of 90% and a power-to-mass ratio of 6 kW/kg. So replacing a 75 kW petrol engine with a 75 kW electric motor saves 85 kg in weight. Sadly, the power to weight ratio of batteries is about 1 kW per kg, so what the electric vehicle gained on the motor, it loses on the batteries.

259 *The bike’s engine uses energy with an efficiency of 0.25.* This and the other assumptions about cycling are confirmed by di Prampero et al. (1979). The drag-area of a cyclist in racing posture is $c_d A = 0.3 \text{ m}^2$. The rolling resistance of a cyclist on a high-quality racing cycle (total weight 73 kg) is 3.2 N.

260 *Figure A.12.*

Prius data from B. Z. Wilson [home.hiwaay.net/~bzwilson/prius/]. BMW data from Phil C. Stuart [www.randomuseless.info/318ti/economy.html].

Further reading: Gabrielli and von Kármán (1950).




Figure A.15. The Wartsila-Sulzer RTA96-C 14-cylinder two-stroke diesel engine. 27 m long and 13.5 m high. www.wartsila.com


B Wind II

The physics of wind power


To estimate the energy in wind, let’s imagine holding up a hoop with area A , facing the wind whose speed is v . Consider the mass of air that passes through that hoop in one second. Here’s a picture of that mass of air just before it passes through the hoop:



And here’s a picture of the same mass of air one second later:



The mass of this piece of air is the product of its density ρ , its area A , and its length, which is v times t , where t is one second.



The kinetic energy of this piece of air is

$$\frac{1}{2}mv^2 = \frac{1}{2}\rho Avt v^2 = \frac{1}{2}\rho Atv^3. \tag{B.1}$$

So the power of the wind, for an area A – that is, the kinetic energy passing across that area per unit time – is

$$\frac{\frac{1}{2}mv^2}{t} = \frac{1}{2}\rho Av^3. \tag{B.2}$$

This formula may look familiar – we derived an identical expression on p255 when we were discussing the power requirement of a moving car.

What’s a typical wind speed? On a windy day, a cyclist really notices the wind direction; if the wind is behind you, you can go much faster than



I’m using this formula again:
mass = density × volume

miles/ hour	km/h	m/s	Beaufort scale
2.2	3.6	1	force 1
7	11	3	force 2
11	18	5	force 3
13	21	6	force 4
16	25	7	
22	36	10	force 5
29	47	13	force 6
36	58	16	force 7
42	68	19	force 8
49	79	22	force 9
60	97	27	force 10
69	112	31	force 11
78	126	35	force 12

Figure B.1. Speeds.




Figure B.2. Flow of air past a windmill. The air is slowed down and splayed out by the windmill.

normal; the speed of such a wind is therefore comparable to the typical speed of the cyclist, which is, let's say, 21 km per hour (13 miles per hour, or 6 metres per second). In Cambridge, the wind is only occasionally this big. Nevertheless, let's use this as a typical British figure (and bear in mind that we may need to revise our estimates).

The density of air is about 1.3 kg per m³. (I usually round this to 1 kg per m³, which is easier to remember, although I haven't done so here.) Then the typical power of the wind per square metre of hoop is

$$\frac{1}{2}\rho v^3 = \frac{1}{2}1.3 \text{ kg/m}^3 \times (6 \text{ m/s})^3 = 140 \text{ W/m}^2. \quad (\text{B.3})$$

Not all of this energy can be extracted by a windmill. The windmill slows the air down quite a lot, but it has to leave the air with *some* kinetic energy, otherwise that slowed-down air would get in the way. Figure B.2 is a cartoon of the actual flow past a windmill. The maximum fraction of the incoming energy that can be extracted by a disc-like windmill was worked out by a German physicist called Albert Betz in 1919. If the departing wind speed is one third of the arriving wind speed, the power extracted is 16/27 of the total power in the wind. 16/27 is 0.59. In practice let's guess that a windmill might be 50% efficient. In fact, real windmills are designed with particular wind speeds in mind; if the wind speed is significantly greater than the turbine's ideal speed, it has to be switched off.

As an example, let's assume a diameter of $d = 25 \text{ m}$, and a hub height of 32 m, which is roughly the size of the lone windmill above the city of Wellington, New Zealand (figure B.3). The power of a single windmill is

$$\begin{aligned} & \text{efficiency factor} \times \text{power per unit area} \times \text{area} \\ &= 50\% \times \frac{1}{2}\rho v^3 \times \frac{\pi}{4}d^2 \end{aligned} \quad (\text{B.4})$$

$$= 50\% \times 140 \text{ W/m}^2 \times \frac{\pi}{4}(25 \text{ m})^2 \quad (\text{B.5})$$

$$= 34 \text{ kW}. \quad (\text{B.6})$$

Indeed, when I visited this windmill on a very breezy day, its meter showed it was generating 60 kW.

To estimate how much power we can get from wind, we need to decide how big our windmills are going to be, and how close together we can pack them.



Figure B.3. The Brooklyn windmill above Wellington, New Zealand, with people providing a scale at the base. On a breezy day, this windmill was producing 60 kW, (1400 kWh per day). Photo by Philip Banks.

How densely could such windmills be packed? Too close and the upwind ones will cast wind-shadows on the downwind ones. Experts say that windmills can't be spaced closer than 5 times their diameter without losing significant power. At this spacing, the power that windmills can generate per unit land area is

$$\frac{\text{power per windmill (B.4)}}{\text{land area per windmill}} = \frac{\frac{1}{2}\rho v^3 \frac{\pi}{8} d^2}{(5d)^2} \tag{B.7}$$

$$= \frac{\pi}{200} \frac{1}{2} \rho v^3 \tag{B.8}$$

$$= 0.016 \times 140 \text{ W/m}^2 \tag{B.9}$$

$$= 2.2 \text{ W/m}^2. \tag{B.10}$$

This number is worth remembering: a wind farm with a wind speed of 6 m/s produces a power of 2 W per m² of land area. Notice that our answer does not depend on the diameter of the windmill. The *ds* cancelled because bigger windmills have to be spaced further apart. Bigger windmills might be a good idea in order to catch bigger windspeeds that exist higher up (the taller a windmill is, the bigger the wind speed it encounters), or because of economies of scale, but those are the only reasons for preferring big windmills.

This calculation depended sensitively on our estimate of the windspeed. Is 6 m/s plausible as a long-term typical windspeed in windy parts of Britain? Figures 4.1 and 4.2 showed windspeeds in Cambridge and Cairngorm. Figure B.6 shows the mean winter and summer windspeeds in eight more locations around Britain. I fear 6 m/s was an overestimate of the typical speed in most of Britain! If we replace 6 m/s by Bedford's




Figure B.4. Wind farm layout.

POWER PER UNIT AREA	
wind farm	2 W/m ²
(speed 6 m/s)	

Table B.5. Facts worth remembering: wind farms.




Figure B.6. Average summer windspeed (dark bar) and average winter windspeed (light bar) in eight locations around Britain. Speeds were measured at the standard weatherman's height of 10 metres. Averages are over the period 1971–2000.

4 m/s as our estimated windspeed, we must scale our estimate down, multiplying it by $(4/6)^3 \simeq 0.3$. (Remember, wind power scales as wind-speed cubed.)

On the other hand, to estimate the typical power, we shouldn't take the mean wind speed and cube it; rather, we should find the mean cube of the windspeed. The average of the cube is bigger than the cube of the average. But if we start getting into these details, things get even more complicated, because real wind turbines don't actually deliver a power proportional to wind-speed cubed. Rather, they typically have just a range of wind-speeds within which they deliver the ideal power; at higher or lower speeds real wind turbines deliver less than the ideal power.

Variation of wind speed with height

Taller windmills see higher wind speeds. The way that wind speed increases with height is complicated and depends on the roughness of the surrounding terrain and on the time of day. As a ballpark figure, doubling the height typically increases wind-speed by 10% and thus increases the power of the wind by 30%.

Some standard formulae for speed v as a function of height z are:

1. According to the wind shear formula from NREL [ydt7uk], the speed varies as a power of the height:

$$v(z) = v_{10} \left(\frac{z}{10 \text{ m}} \right)^\alpha,$$

where v_{10} is the speed at 10 m, and a typical value of the exponent α is 0.143 or $1/7$. The one-seventh law ($v(z)$ is proportional to $z^{1/7}$) is used by Elliott et al. (1991), for example.

2. The wind shear formula from the Danish Wind Industry Association [yaoonz] is

$$v(z) = v_{\text{ref}} \frac{\log(z/z_0)}{\log(z_{\text{ref}}/z_0)},$$

where z_0 is a parameter called the roughness length, and v_{ref} is the speed at a reference height z_{ref} such as 10 m. The roughness length for typical countryside (agricultural land with some houses and sheltering hedgerows with some 500-m intervals – “roughness class 2”) is $z_0 = 0.1 \text{ m}$.

In practice, these two wind shear formulae give similar numerical answers. That's not to say that they are accurate at all times however. Van den Berg (2004) suggests that different wind profiles often hold at night.




Figure B.7. Top: Two models of wind speed and wind power as a function of height. DWIA = Danish Wind Industry Association; NREL = National Renewable Energy Laboratory. For each model the speed at 10 m has been fixed to 6 m/s. For the Danish Wind model, the roughness length is set to $z_0 = 0.1 \text{ m}$. Bottom: The power density (the power per unit of upright area) according to each of these models.



Figure B.8. The qr5 from quietrevolution.co.uk. Not a typical windmill.

Standard windmill properties

The typical windmill of today has a rotor diameter of around 54 metres centred at a height of 80 metres; such a machine has a “capacity” of 1 MW. The “capacity” or “peak power” is the *maximum* power the windmill can generate in optimal conditions. Usually, wind turbines are designed to start running at wind speeds somewhere around 3 to 5 m/s and to stop if the wind speed reaches gale speeds of 25 m/s. The actual average power delivered is the “capacity” multiplied by a factor that describes the fraction of the time that wind conditions are near optimal. This factor, sometimes called the “load factor” or “capacity factor,” depends on the site; a typical load factor for a *good* site in the UK is 30%. In the Netherlands, the typical load factor is 22%; in Germany, it is 19%.

Other people’s estimates of wind farm power per unit area

In the government’s study [www.world-nuclear.org/policy/DTI-PIU.pdf] the UK onshore wind resource is estimated using an assumed wind farm power per unit area of at most 9 W/m^2 (capacity, not average production). If the capacity factor is 33% then the average power production would be 3 W/m^2 .

The London Array is an offshore wind farm planned for the outer Thames Estuary. With its 1 GW capacity, it is expected to become the world’s largest offshore wind farm. The completed wind farm will consist of 271 wind turbines in 245 km^2 [6086ec] and will deliver an average power of 3100 GWh per year (350 MW). (Cost £1.5 bn.) That’s a power per unit area of $350 \text{ MW}/245 \text{ km}^2 = 1.4 \text{ W/m}^2$. This is lower than other offshore farms because, I guess, the site includes a big channel (Knock Deep) that’s too deep (about 20 m) for economical planting of turbines.

I’m more worried about what these plans [for the proposed London Array wind farm] will do to this landscape and our way of life than I ever was about a Nazi invasion on the beach.

Bill Boggia of Graveney, where the undersea cables
of the wind farm will come ashore.

Queries

What about micro-generation? If you plop one of those mini-turbines on your roof, what energy can you expect it to deliver?

Assuming a windspeed of 6 m/s, which, as I said before, is *above* the average for most parts of Britain; and assuming a diameter of 1 m, the power delivered would be 50 W. That's 1.3 kWh per day – not very much. And in reality, in a typical urban location in England, a microturbine delivers just 0.2 kWh per day – see p66.

Perhaps the worst windmills in the world are a set in Tsukuba City, Japan, which actually consume more power than they generate. Their installers were so embarrassed by the stationary turbines that they imported power to make them spin so that they looked like they were working! [6bkvbn]

Notes and further reading

page no.

- 264 *The maximum fraction of the incoming energy that can be extracted by a disc-like windmill...* There is a nice explanation of this on the Danish Wind Industry Association's website. [yekdaa].
- 267 *Usually, wind turbines are designed to start running at wind speeds around 3 to 5 m/s.* [ymfbsn].
- *a typical load factor for a good site is 30%.* In 2005, the average load factor of all major UK wind farms was 28% [ypvbvd]. The load factor varied during the year, with a low of 17% in June and July. The load factor for the best region in the country – Caithness, Orkney and the Shetlands – was 33%. The load factors of the two offshore wind farms operating in 2005 were 36% for North Hoyle (off North Wales) and 29% for Scroby Sands (off Great Yarmouth). Average load factors in 2006 for ten regions were: Cornwall 25%; Mid-Wales 27%; Cambridgeshire and Norfolk 25%; Cumbria 25%; Durham 16%; Southern Scotland 28%; Orkney and Shetlands 35%; Northeast Scotland 26%; Northern Ireland 31%; offshore 29%. [wbd8o]

Watson et al. (2002) say a minimum annual mean wind speed of 7.0 m/s is currently thought to be necessary for commercial viability of wind power. About 33% of UK land area has such speeds.



Figure B.9. An Ampair “600 W” micro-turbine. The average power generated by this micro-turbine in Leamington Spa is 0.037 kWh per day (1.5 W).



Figure B.10. A 5.5-m diameter Iskra 5 kW turbine [www.iskrawind.com] having its annual check-up. This turbine, located in Hertfordshire (not the windiest of locations in Britain), mounted at a height of 12 m, has an average output of 11 kWh per day. A wind farm of machines with this performance, one per 30 m × 30 m square, would have a power per unit area of 0.5 W/m².

C Planes II

What we need to do is to look at how you make air travel more energy efficient, how you develop the new fuels that will allow us to burn less energy and emit less.

Tony Blair

Hoping for the best is not a policy, it is a delusion.

Emily Armistead, Greenpeace

What are the fundamental limits of travel by flying? Does the physics of flight require an unavoidable use of a certain amount of energy, per ton, per kilometre flown? What's the maximum distance a 300-ton Boeing 747 can fly? What about a 1-kg bar-tailed godwit or a 100-gram Arctic tern?

Just as Chapter 3, in which we estimated consumption by cars, was followed by Chapter A, offering a model of where the energy goes in cars, this chapter fills out Chapter 5, discussing where the energy goes in planes. The only physics required is Newton's laws of motion, which I'll describe when they're needed.

This discussion will allow us to answer questions such as "would air travel consume much less energy if we travelled in slower propellor-driven planes?" There's a lot of equations ahead: I hope you enjoy them!

How to fly

Planes (and birds) move through air, so, just like cars and trains, they experience a drag force, and much of the energy guzzled by a plane goes into pushing the plane along against this force. Additionally, unlike cars and trains, planes have to expend energy *in order to stay up*.

Planes stay up by throwing air down. When the plane pushes down on air, the air pushes up on the plane (because Newton's third law tells it to). As long as this upward push, which is called lift, is big enough to balance the downward weight of the plane, the plane avoids plummeting downwards.

When the plane throws air down, it gives that air kinetic energy. So creating lift requires energy. The total power required by the plane is the sum of the power required to create lift and the power required to overcome ordinary drag. (The power required to create lift is usually called "induced drag," by the way. But I'll call it the lift power, P_{lift} .)

The two equations we'll need, in order to work out a theory of flight, are Newton's second law:

$$\text{force} = \text{rate of change of momentum}, \quad (\text{C.1})$$



Figure C.1. Birds: two Arctic terns, a bar-tailed godwit, and a Boeing 747.




Figure C.2. A plane encounters a stationary tube of air. Once the plane has passed by, the air has been thrown downwards by the plane. The force exerted by the plane on the air to accelerate it downwards is equal and opposite to the upwards force exerted on the plane by the air.




Figure C.3. Our cartoon assumes that the plane leaves a sausage of air moving down in its wake. A realistic picture involves a more complex swirling flow. For the real thing, see figure C.4.

and Newton's third law, which I just mentioned:

$$\text{force exerted on A by B} = - \text{force exerted on B by A.} \quad (\text{C.2})$$

If you don't like equations, I can tell you the punchline now: we're going to find that the power required to create lift turns out to be *equal* to the power required to overcome drag. So the requirement to "stay up" *doubles* the power required.

Let's make a cartoon of the lift force on a plane moving at speed v . In a time t the plane moves a distance vt and leaves behind it a sausage of downward-moving air (figure C.2). We'll call the cross-sectional area of this sausage A_s . This sausage's diameter is roughly equal to the wingspan w of the plane. (Within this large sausage is a smaller sausage of swirling turbulent air with cross-sectional area similar to the frontal area of the plane's body.) Actually, the details of the air flow are much more interesting than this sausage picture: each wing tip leaves behind it a vortex, with the air between the wingtips moving down fast, and the air beyond (outside) the wingtips moving up (figures C.3 & C.4). This upward-moving air is exploited by birds flying in formation: just behind the tip of a bird's wing is a sweet little updraft. Anyway, let's get back to our sausage.

The sausage's mass is

$$m_{\text{sausage}} = \text{density} \times \text{volume} = \rho vt A_s. \quad (\text{C.3})$$

Let's say the whole sausage is moving down with speed u , and figure out what u needs to be in order for the plane to experience a lift force equal to



Figure C.4. Air flow behind a plane. Photo by NASA Langley Research Center.

its weight mg . The downward momentum of the sausage created in time t is

$$\text{mass} \times \text{velocity} = m_{\text{sausage}}u = \rho vt A_s u. \quad (\text{C.4})$$

And by Newton's laws this must equal the momentum delivered by the plane's weight in time t , namely,

$$mgt. \quad (\text{C.5})$$

Rearranging this equation,

$$\rho vt A_s u = mgt, \quad (\text{C.6})$$

we can solve for the required downward sausage speed,

$$u = \frac{mg}{\rho v A_s}.$$

Interesting! The sausage speed is *inversely* related to the plane's speed v . A slow-moving plane has to throw down air harder than a fast-moving plane, because it encounters less air per unit time. That's why landing planes, travelling slowly, have to extend their flaps: so as to create a larger and steeper wing that deflects air more.

What's the energetic cost of pushing the sausage down at the required speed u ? The power required is

$$P_{\text{lift}} = \frac{\text{kinetic energy of sausage}}{\text{time}} \quad (\text{C.7})$$

$$= \frac{1}{t} \frac{1}{2} m_{\text{sausage}} u^2 \quad (\text{C.8})$$

$$= \frac{1}{2t} \rho vt A_s \left(\frac{mg}{\rho v A_s} \right)^2 \quad (\text{C.9})$$

$$= \frac{1}{2} \frac{(mg)^2}{\rho v A_s}. \quad (\text{C.10})$$

The total power required to keep the plane going is the sum of the drag power and the lift power:

$$P_{\text{total}} = P_{\text{drag}} + P_{\text{lift}} \quad (\text{C.11})$$

$$= \frac{1}{2} c_d \rho A_p v^3 + \frac{1}{2} \frac{(mg)^2}{\rho v A_s}, \quad (\text{C.12})$$

where A_p is the frontal area of the plane and c_d is its drag coefficient (as in Chapter A).

The fuel-efficiency of the plane, expressed as the energy per distance travelled, would be

$$\frac{\text{energy}}{\text{distance}} \Big|_{\text{ideal}} = \frac{P_{\text{total}}}{v} = \frac{1}{2} c_d \rho A_p v^2 + \frac{1}{2} \frac{(mg)^2}{\rho v^2 A_s}, \quad (\text{C.13})$$

if the plane turned its fuel's power into drag power and lift power perfectly efficiently. (Incidentally, another name for "energy per distance travelled" is "force," and we can recognize the two terms above as the drag force $\frac{1}{2}c_d\rho A_p v^2$ and the lift-related force $\frac{1}{2}\frac{(mg)^2}{\rho v^2 A_s}$. The sum is the force, or "thrust," that specifies exactly how hard the engines have to push.)

Real jet engines have an efficiency of about $\epsilon = 1/3$, so the energy-per-distance of a plane travelling at speed v is

$$\frac{\text{energy}}{\text{distance}} = \frac{1}{\epsilon} \left(\frac{1}{2}c_d\rho A_p v^2 + \frac{1}{2}\frac{(mg)^2}{\rho v^2 A_s} \right). \quad (\text{C.14})$$

This energy-per-distance is fairly complicated; but it simplifies greatly if we assume that the plane is *designed* to fly at the speed that *minimizes* the energy-per-distance. The energy-per-distance, you see, has got a sweet-spot as a function of v (figure C.5). The sum of the two quantities $\frac{1}{2}c_d\rho A_p v^2$ and $\frac{1}{2}\frac{(mg)^2}{\rho v^2 A_s}$ is smallest when the two quantities are equal. This phenomenon is delightfully common in physics and engineering: two things that don't obviously *have* to be equal *are* actually equal, or equal within a factor of 2.

So, this equality principle tells us that the optimum speed for the plane is such that

$$c_d\rho A_p v^2 = \frac{(mg)^2}{\rho v^2 A_s}, \quad (\text{C.15})$$

i.e.,

$$\rho v_{\text{opt}}^2 = \frac{mg}{\sqrt{c_d A_p A_s}}, \quad (\text{C.16})$$

This defines the optimum speed if our cartoon of flight is accurate; the cartoon breaks down if the engine efficiency ϵ depends significantly on speed, or if the speed of the plane exceeds the speed of sound (330 m/s); above the speed of sound, we would need a different model of drag and lift.

Let's check our model by seeing what it predicts is the optimum speed for a 747 and for an albatross. We must take care to use the correct air-density: if we want to estimate the optimum cruising speed for a 747 at 30 000 feet, we must remember that air density drops with increasing altitude z as $\exp(-mgz/kT)$, where m is the mass of nitrogen or oxygen molecules, and kT is the thermal energy (Boltzmann's constant times absolute temperature). The density is about 3 times smaller at that altitude.

The predicted optimal speeds (table C.6) are more accurate than we have a right to expect! The 747's optimal speed is predicted to be 540 mph, and the albatross's, 32 mph – both very close to the true cruising speeds of the two birds (560 mph and 30–55 mph respectively).

Let's explore a few more predictions of our cartoon. We can check whether the force (C.13) is compatible with the known thrust of the 747. Remembering that at the optimal speed, the two forces are equal, we just




Figure C.5. The force required to keep a plane moving, as a function of its speed v , is the sum of an ordinary drag force $\frac{1}{2}c_d\rho A_p v^2$ – which increases with speed – and the lift-related force (also known as the induced drag) $\frac{1}{2}\frac{(mg)^2}{\rho v^2 A_s}$ – which decreases with speed. There is an ideal speed, v_{optimal} , at which the force required is minimized. The force is an energy per distance, so minimizing the force also minimizes the fuel per distance. To optimize the fuel efficiency, fly at v_{optimal} . This graph shows our cartoon's estimate of the thrust required, in kilonewtons, for a Boeing 747 of mass 319 t, wingspan 64.4 m, drag coefficient 0.03, and frontal area 180 m², travelling in air of density $\rho = 0.41 \text{ kg/m}^3$ (the density at a height of 10 km), as a function of its speed v in m/s. Our model has an optimal speed $v_{\text{optimal}} = 220 \text{ m/s}$ (540 mph). For a cartoon based on sausages, this is a good match to real life!

BIRD		747	Albatross
Designer		Boeing	natural selection
Mass (fully-laden)	m	363 000 kg	8 kg
Wingspan	w	64.4 m	3.3 m
Area*	A_p	180 m ²	0.09 m ²
Density	ρ	0.4 kg/m ³	1.2 kg/m ³
Drag coefficient	c_d	0.03	0.1
Optimum speed	v_{opt}	220 m/s = 540 mph	14 m/s = 32 mph

Table C.6. Estimating the optimal speeds for a jumbo jet and an albatross.
 * Frontal area estimated for 747 by taking cabin width (6.1 m) times estimated height of body (10 m) and adding double to allow for the frontal area of engines, wings, and tail; for albatross, frontal area of 1 square foot estimated from a photograph.

need to pick one of them and double it:

$$\text{force} = \frac{\text{energy}}{\text{distance}} \Big|_{\text{ideal}} = \frac{1}{2} c_d \rho A_p v^2 + \frac{1}{2} \frac{(mg)^2}{\rho v^2 A_s} \quad (\text{C.17})$$

$$= c_d \rho A_p v_{\text{opt}}^2 \quad (\text{C.18})$$

$$= c_d \rho A_p \frac{mg}{\rho (c_d A_p A_s)^{1/2}} \quad (\text{C.19})$$

$$= \left(\frac{c_d A_p}{A_s} \right)^{1/2} mg. \quad (\text{C.20})$$

Let's define the filling factor f_A to be the area ratio:

$$f_A = \frac{A_p}{A_s}. \quad (\text{C.21})$$

(Think of f_A as the fraction of the square occupied by the plane in figure C.7.) Then

$$\text{force} = (c_d f_A)^{1/2} (mg). \quad (\text{C.22})$$

Interesting! Independent of the density of the fluid through which the plane flies, the required thrust (for a plane travelling at the optimal speed) is just a dimensionless constant $(c_d f_A)^{1/2}$ times the weight of the plane. This constant, by the way, is known as the drag-to-lift ratio of the plane. (The lift-to-drag ratio has a few other names: the glide number, glide ratio, aerodynamic efficiency, or finesse; typical values are shown in table C.8.)

Taking the jumbo jet's figures, $c_d \simeq 0.03$ and $f_A \simeq 0.04$, we find the required thrust is

$$(c_d f_A)^{1/2} mg = 0.036 mg = 130 \text{ kN}. \quad (\text{C.23})$$

How does this agree with the 747's spec sheets? In fact each of the 4 engines has a maximum thrust of about 250 kN, but this maximum thrust is used only during take-off. During cruise, the thrust is much smaller:




Figure C.7. Frontal view of a Boeing 747, used to estimate the frontal area A_p of the plane. The square has area A_s (the square of the wingspan).

Airbus A320	17
Boeing 767-200	19
Boeing 747-100	18
Common Tern	12
Albatross	20

Table C.8. Lift-to-drag ratios.

the thrust of a cruising 747 is 200 kN, just 50% more than our cartoon suggested. Our cartoon is a little bit off because our estimate of the drag-to-lift ratio was a little bit low.

This thrust can be used directly to deduce the transport efficiency achieved by any plane. We can work out two sorts of transport efficiency: the energy cost of moving *weight* around, measured in kWh per ton-kilometre; and the energy cost of moving people, measured in kWh per 100 passenger-kilometres.

Efficiency in weight terms

Thrust is a force, and a force is an energy per unit distance. The total energy used per unit distance is bigger by a factor $(1/\epsilon)$, where ϵ is the efficiency of the engine, which we'll take to be $1/3$.

Here's the gross transport cost, defined to be the energy per unit weight (of the entire craft) per unit distance:

$$\text{transport cost} = \frac{1}{\epsilon} \frac{\text{force}}{\text{mass}} \quad (\text{C.24})$$

$$= \frac{1}{\epsilon} \frac{(c_d f_A)^{1/2} m g}{m} \quad (\text{C.25})$$

$$= \frac{(c_d f_A)^{1/2}}{\epsilon} g. \quad (\text{C.26})$$

So the transport cost is just a dimensionless quantity (related to a plane's shape and its engine's efficiency), multiplied by g , the acceleration due to gravity. Notice that this gross transport cost applies to all planes, but depends only on three simple properties of the plane: its drag coefficient, the shape of the plane, and its engine efficiency. It doesn't depend on the size of the plane, nor on its weight, nor on the density of air. If we plug in $\epsilon = 1/3$ and assume a lift-to-drag ratio of 20 we find the gross transport cost of *any* plane, according to our cartoon, is

$$0.15 g$$

or

$$0.4 \text{ kWh/ton-km.}$$

Can planes be improved?

If engine efficiency can be boosted only a tiny bit by technological progress, and if the shape of the plane has already been essentially perfected, then there is little that can be done about the dimensionless quantity. The transport efficiency is close to its physical limit. The aerodynamics community say that the shape of planes could be improved a little by a switch to blended-wing bodies, and that the drag coefficient could be reduced a



Figure C.9. Cessna 310N: 60 kWh per 100 passenger-km. A Cessna 310 Turbo carries 6 passengers (including 1 pilot) at a speed of 370 km/h. Photograph by Adrian Pingstone.

little by laminar flow control, a technology that reduces the growth of turbulence over a wing by sucking a little air through small perforations in the surface (Braslow, 1999). Adding laminar flow control to existing planes would deliver a 15% improvement in drag coefficient, and the change of shape to blended-wing bodies is predicted to improve the drag coefficient by about 18% (Green, 2006). And equation (C.26) says that the transport cost is proportional to the square root of the drag coefficient, so improvements of c_d by 15% or 18% would improve transport cost by 7.5% and 9% respectively.

This gross transport cost is the energy cost of moving weight around, *including the weight of the plane itself*. To estimate the energy required to move freight by plane, per unit weight of freight, we need to divide by the fraction that is cargo. For example, if a full 747 freighter is about 1/3 cargo, then its transport cost is

$$0.45 g,$$

or roughly 1.2 kWh/ton-km. This is just a little bigger than the transport cost of a truck, which is 1 kWh/ton-km.

Transport efficiency in terms of bodies

Similarly, we can estimate a passenger transport-efficiency for a 747.

$$\begin{aligned} & \text{transport efficiency (passenger-km per litre of fuel)} \\ &= \text{number of passengers} \times \frac{\text{energy per litre}}{\text{thrust}} \quad (\text{C.27}) \end{aligned}$$

$$= \text{number of passengers} \times \frac{\epsilon \times \text{energy per litre}}{\text{thrust}} \quad (\text{C.28})$$

$$= 400 \times \frac{1}{3} \frac{38 \text{ MJ/litre}}{200\,000 \text{ N}} \quad (\text{C.29})$$

$$= 25 \text{ passenger-km per litre} \quad (\text{C.30})$$

This is a bit more efficient than a typical single-occupant car (12 km per litre). So travelling by plane is more energy-efficient than car if there are only one or two people in the car; and cars are more efficient if there are three or more passengers in the vehicle.

Key points

We've covered quite a lot of ground! Let's recap the key ideas. Half of the work done by a plane goes into *staying up*; the other half goes into *keeping going*. The fuel efficiency at the optimal speed, expressed as an energy-per-distance-travelled, was found in the force (C.22), and it was simply proportional to the weight of the plane; the constant of proportionality is the drag-to-lift ratio, which is determined by the shape of the plane.



Figure C.10. “Fasten your cufflinks.” A Bombardier Learjet 60XR carrying 8 passengers at 780 km/h has a transport cost of **150 kWh per 100 passenger-km**. Photograph by Adrian Pingstone.

So whereas lowering speed-limits for cars would reduce the energy consumed per distance travelled, there is no point in considering speed-limits for planes. Planes that are up in the air have optimal speeds, different for each plane, depending on its weight, and they already go at their optimal speeds. If you ordered a plane to go slower, its energy consumption would *increase*. The only way to make a plane consume fuel more efficiently is to put it on the ground and stop it. Planes have been fantastically optimized, and there is no prospect of significant improvements in plane efficiency. (See pages 37 and 132 for further discussion of the notion that new super-jumbos are “far more efficient” than old jumbos; and p35 for discussion of the notion that turboprops are “far more efficient” than jets.)



Figure C.11. Boeing 737-700: 30 kWh per 100 passenger-km. Photograph © Tom Collins.

Range

Another prediction we can make is, what’s the range of a plane or bird – the biggest distance it can go without refuelling? You might think that bigger planes have a bigger range, but the prediction of our model is startlingly simple. The range of the plane, the maximum distance it can go before refuelling, is proportional to its velocity and to the total energy of the fuel, and inversely proportional to the rate at which it guzzles fuel:

$$\text{range} = v_{\text{opt}} \frac{\text{energy}}{\text{power}} = \frac{\text{energy} \times \epsilon}{\text{force}}. \quad (\text{C.31})$$

Now, the total energy of fuel is the calorific value of the fuel, C (in joules per kilogram), times its mass; and the mass of fuel is some fraction f_{fuel} of the total mass of the plane. So

$$\text{range} = \frac{\text{energy} \epsilon}{\text{force}} = \frac{C m \epsilon f_{\text{fuel}}}{(c_d f_A)^{1/2} (m g)} = \frac{\epsilon f_{\text{fuel}}}{(c_d f_A)^{1/2}} \frac{C}{g}. \quad (\text{C.32})$$

It’s hard to imagine a simpler prediction: the range of any bird or plane is the product of a dimensionless factor $\left(\frac{\epsilon f_{\text{fuel}}}{(c_d f_A)^{1/2}} \right)$ which takes into account the engine efficiency, the drag coefficient, and the bird’s geometry, with a fundamental distance,

$$\frac{C}{g},$$

which is a property of the fuel and gravity, and nothing else. No bird size, no bird mass, no bird length, no bird width; no dependence on the fluid density.

So what is this magic length? It’s the same distance whether the fuel is goose fat or jet fuel: both these fuels are essentially hydrocarbons $(\text{CH}_2)_n$. Jet fuel has a calorific value of $C = 40 \text{ MJ per kg}$. The distance associated with jet fuel is

$$d_{\text{Fuel}} = \frac{C}{g} = 4000 \text{ km}. \quad (\text{C.33})$$

The range of the bird is the intrinsic range of the fuel, 4000 km, times a factor $\left(\frac{\epsilon f_{\text{fuel}}}{(c_d f_A)^{1/2}}\right)$. If our bird has engine efficiency $\epsilon = 1/3$ and drag-to-lift ratio $(c_d f_A)^{1/2} \simeq 1/20$, and if nearly half of the bird is fuel (a fully-laden 747 is 46% fuel), we find that all birds and planes, of whatever size, have the same range: about three times the fuel's distance – roughly 13 000 km.

This figure is again close to the true answer: the nonstop flight record for a 747 (set on March 23–24, 1989) was a distance of 16 560 km.

And the claim that the range is independent of bird size is supported by the observation that birds of all sizes, from great geese down to dainty swallows and arctic tern migrate intercontinental distances. The longest recorded non-stop flight by a bird was a distance of 11 000 km, by a bar-tailed godwit.

How far did Steve Fossett go in the specially-designed Scaled Composites Model 311 Virgin Atlantic GlobalFlyer? 41 467 km. [33ptcg] An unusual plane: 83% of its take-off weight was fuel; the flight made careful use of the jet-stream to boost its distance. Fragile, the plane had several failures along the way.

One interesting point brought out by this cartoon: if we ask “what's the optimum air-density to fly in?”, we find that the *thrust* required (C.20) at the optimum speed is independent of the density. So our cartoon plane would be equally happy to fly at any height; there isn't an optimum density; the plane could achieve the same miles-per-gallon in any density; but the optimum *speed* does depend on the density ($v^2 \sim 1/\rho$, equation (C.16)). So all else being equal, our cartoon plane would have the shortest journey time if it flew in the lowest-density air possible. Now real engines' efficiencies aren't independent of speed and air density. As a plane gets lighter by burning fuel, our cartoon says its optimal speed at a given density would reduce ($v^2 \sim mg/(\rho(c_d A_p A_s)^{1/2})$). So a plane travelling in air of constant density should slow down a little as it gets lighter. But a plane can both keep going at a *constant speed* and continue flying at its *optimal* speed if it increases its altitude so as to reduce the air density. Next time you're on a long-distance flight, you could check whether the pilot increases the cruising height from, say, 31 000 feet to 39 000 feet by the end of the flight.

You can think of d_{Fuel} as the distance that the fuel could throw itself if it suddenly converted all its chemical energy to kinetic energy and launched itself on a parabolic trajectory with no air resistance. [To be precise, the distance achieved by the optimal parabola is twice C/g .] This distance is also the *vertical* height to which the fuel could throw itself if there were no air resistance. Another amusing thing to notice is that the calorific value of a fuel C , which I gave in joules per kilogram, is also a squared-velocity (just as the energy-to-mass ratio E/m in Einstein's $E = mc^2$ is a squared-velocity, c^2): 40×10^6 J per kg is $(6000 \text{ m/s})^2$. So one way to think about fat is “fat is 6000 metres per second.” If you want to lose weight by going jogging, 6000 m/s (12 000 mph) is the speed you should aim for in order to lose it all in one giant leap.

How would a hydrogen plane perform?

We've already argued that the efficiency of flight, in terms of energy per ton-km, is just a simple dimensionless number times g . Changing the fuel isn't going to change this fundamental argument. Hydrogen-powered planes are worth discussing if we're hoping to reduce climate-changing emissions. They might also have better range. But don't expect them to be radically more energy-efficient.

Possible areas for improvement of plane efficiency

Formation flying in the style of geese could give a 10% improvement in fuel efficiency (because the lift-to-drag ratio of the formation is higher than that of a single aircraft), but this trick relies, of course, on the geese wanting to migrate to the same destination at the same time.

Optimizing the hop lengths: long-range planes (designed for a range of say 15 000 km) are not quite as fuel-efficient as shorter-range planes, because they have to carry extra fuel, which makes less space for cargo and passengers. It would be more energy-efficient to fly shorter hops in shorter-range planes. The sweet spot is when the hops are about 5000 km long, so typical long-distance journeys would have one or two refuelling stops (Green, 2006). Multi-stage long-distance flying might be about 15% more fuel-efficient; but of course it would introduce other costs.

Eco-friendly aeroplanes

Occasionally you may hear about people making eco-friendly aeroplanes. Earlier in this chapter, however, our cartoon made the assertion that the transport cost of *any* plane is about

$$0.4 \text{ kWh/ton-km.}$$

According to the cartoon, the only ways in which a plane could significantly improve on this figure are to reduce air resistance (perhaps by some new-fangled vacuum-cleaners-in-the-wings trick) or to change the geometry of the plane (making it look more like a glider, with immensely wide wings compared to the fuselage, or getting rid of the fuselage altogether).

So, let's look at the latest news story about "eco-friendly aviation" and see whether one of these planes can beat the 0.4 kWh per ton-km benchmark. If a plane uses less than 0.4 kWh per ton-km, we might conclude that the cartoon is defective.

The Electra, a wood-and-fabric single-seater, flew for 48 minutes for 50 km around the southern Alps [6r32hf]. The Electra has a 9-m wingspan and an 18-kW electric motor powered by 48 kg of lithium-polymer batteries. The aircraft's take-off weight is 265 kg (134 kg of aircraft, 47 kg of batteries, and 84 kg of human cargo). On 23rd December, 2007 it flew a distance of 50 km. If we assume that the battery's energy density was 130 Wh/kg, and that the flight used 90% of a full charge (5.5 kWh), the transport cost was roughly

$$0.4 \text{ kWh/ton-km,}$$

which exactly matches our cartoon. This electrical plane is not a lower-energy plane than a normal fossil-sucker.

Of course, this doesn't mean that electric planes are not interesting. If one could replace traditional planes by alternatives with equal energy



Figure C.12. The Electra F-WMDJ: 11 kWh per 100 p-km. Photo by Jean-Bernard Gache. www.apame.eu




Figure C.13. Hydrofoil.
Photograph by Georgios Pazios.

consumption but no carbon emissions, that would certainly be a useful technology. And, as a person-transporter, the Electra delivers a respectable 11 kWh per 100 p-km, similar to the electric car in our transport diagram on p128. But in this book the bottom line is always: “where is the energy to come from?”

Many boats are birds too

Some time after writing this cartoon of flight, I realized that it applies to more than just the birds of the air – it applies to hydrofoils, and to other high-speed watercraft too – all those that ride higher in the water when moving.

Figure C.13 shows the principle of the hydrofoil. The weight of the craft is supported by a tilted underwater wing, which may be quite tiny compared with the craft. The wing generates lift by throwing fluid down, just like the plane of figure C.2. If we assume that the drag is dominated by the drag on the wing, and that the wing dimensions and vessel speed have been optimized to minimize the energy expended per unit distance, then the best possible transport cost, in the sense of energy per ton-kilometre, will be just the same as in equation (C.26):

$$\frac{(c_d f_A)^{1/2}}{\epsilon} g, \quad (\text{C.34})$$

where c_d is the drag coefficient of the underwater wing, f_A is the dimensionless area ratio defined before, ϵ is the engine efficiency, and g is the acceleration due to gravity.

Perhaps c_d and f_A are not quite the same as those of an optimized aeroplane. But the remarkable thing about this theory is that it has no dependence on the density of the fluid through which the wing is flying. So our ballpark prediction is that the transport cost (energy-per-distance-per-weight, including the vehicle weight) of a hydrofoil is *the same* as the transport cost of an aeroplane! Namely, roughly 0.4 kWh per ton-km.

For vessels that skim the water surface, such as high-speed catamarans and water-skiers, an accurate cartoon should also include the energy going into making waves, but I'm tempted to guess that this hydrofoil theory is still roughly right.

I've not yet found data on the transport-cost of a hydrofoil, but some data for a passenger-carrying catamaran travelling at 41 km/h seem to agree pretty well: it consumes roughly 1 kWh per ton-km.

It's quite a surprise to me to learn that an island hopper who goes from island to island by plane not only gets there faster than someone who hops by boat – he quite probably uses less energy too.

Other ways of staying up

Airships

This chapter has emphasized that planes can't be made more energy-efficient by slowing them down, because any benefit from reduced air-resistance is more than cancelled by having to chuck air down harder. Can this problem be solved by switching strategy: not throwing air down, but being as light as air instead? An airship, blimp, zeppelin, or dirigible uses an enormous helium-filled balloon, which is lighter than air, to counteract the weight of its little cabin. The disadvantage of this strategy is that the enormous balloon greatly increases the air resistance of the vehicle.

The way to keep the energy cost of an airship (per weight, per distance) low is to move slowly, to be fish-shaped, and to be very large and long. Let's work out a cartoon of the energy required by an idealized airship.

I'll assume the balloon is ellipsoidal, with cross-sectional area A and length L . The volume is $V = \frac{2}{3}AL$. If the airship floats stably in air of density ρ , the total mass of the airship, including its cargo and its helium, must be $m_{\text{total}} = \rho V$. If it moves at speed v , the force of air resistance is

$$F = \frac{1}{2}c_d A \rho v^2, \quad (\text{C.35})$$

where c_d is the drag coefficient, which, based on aeroplanes, we might expect to be about 0.03. The energy expended, per unit distance, is equal to F divided by the efficiency ϵ of the engines. So the gross transport cost – the energy used per unit distance per unit mass – is

$$\frac{F}{\epsilon m_{\text{total}}} = \frac{\frac{1}{2}c_d A \rho v^2}{\epsilon \rho \frac{2}{3}AL} \quad (\text{C.36})$$




Figure C.14. The 239 m-long USS Akron (ZRS-4) flying over Manhattan. It weighed 100 t and could carry 83 t. Its engines had a total power of 3.4 MW, and it could transport 89 personnel and a stack of weapons at 93 km/h. It was also used as an aircraft carrier.




Figure C.15. An ellipsoidal airship.

$$= \frac{3}{4\epsilon} c_d \frac{v^2}{L} \quad (\text{C.37})$$

That's a rather nice result! The gross transport cost of this idealized airship depends only its speed v and length L , not on the density ρ of the air, nor on the airship's frontal area A .

This cartoon also applies without modification to submarines. The gross transport cost (in kWh per ton-km) of an airship is just the same as the gross transport cost of a submarine of identical length and speed. The submarine will contain 1000 times more mass, since water is 1000 times denser than air; and it will cost 1000 times more to move it along. The only difference between the two will be the advertising revenue.

So, let's plug in some numbers. Let's assume we desire to travel at a speed of 80 km/h (so that crossing the Atlantic takes three days). In SI units, that's 22 m/s. Let's assume an efficiency ϵ of 1/4. To get the best possible transport cost, what is the longest blimp we can imagine? The Hindenburg was 245 m long. If we say $L = 400$ m, we find the transport cost is:

$$\frac{F}{\epsilon m_{\text{total}}} = 3 \times 0.03 \frac{(22 \text{ m/s})^2}{400 \text{ m}} = 0.1 \text{ m/s}^2 = \mathbf{0.03 \text{ kWh/t-km}}.$$

If useful cargo made up half of the vessel's mass, the net transport cost of this monster airship would be $\mathbf{0.06 \text{ kWh/t-km}}$ – similar to rail.

Ekranoplans

The ekranoplan, or water-skimming wingship, is a ground-effect aircraft: an aircraft that flies very close to the surface of the water, obtaining its lift not from hurling air down like a plane, nor from hurling water down like a hydrofoil or speed boat, but by sitting on a cushion of compressed air sandwiched between its wings and the nearby surface. You can demonstrate the ground effect by flicking a piece of card across a flat table. Maintaining this air-cushion requires very little energy, so the ground-effect aircraft, in energy terms, is a lot like a surface vehicle with no rolling resistance. Its main energy expenditure is associated with air resistance. Remember that for a plane at its optimal speed, half of its energy expenditure is associated with air resistance, and half with throwing air down.

The Soviet Union developed the ekranoplan as a military transport vehicle and missile launcher in the Khrushchev era. The Lun ekranoplan could travel at 500 km/h, and the total thrust of its eight engines was 1000 kN, though this total was not required once the vessel had risen clear of the water. Assuming the cruising thrust was one quarter of the maximum; that the engines were 30% efficient; and that of its 400-ton weight, 100 tons were cargo, this vehicle had a net freight-transport cost of $\mathbf{2 \text{ kWh per ton-km}}$. I imagine that, if perfected for non-military freight transport, the ekranoplan might have a freight-transport cost about half that of an ordinary aeroplane.




Figure C.16. The Lun ekranoplan – slightly longer and heavier than a Boeing 747. Photographs: A. Belyaev.

Mythconceptions

The plane was going anyway, so my flying was energy-neutral.

This is false for two reasons. First, your extra weight on the plane requires extra energy to be consumed in keeping you up. Second, airlines respond to demand by flying more planes.

Notes and further reading

page no.

- 272 *Boeing 747*. Drag coefficient for 747 from www.aerospaceweb.org. Other 747 data from [2af5gw]. Albatross facts from [32judd].
- *Real jet engines have an efficiency of about $\epsilon = 1/3$* . Typical engine efficiencies are in the range 23%–36% [adg.stanford.edu/aa241/propulsion/sfc.html]. For typical aircraft, overall engine efficiency ranges between 20% and 40%, with the best bypass engines delivering 30–37% when cruising [www.grida.no/climate/ipcc/aviation/097.htm]. You can't simply pick the most efficient engine however, since it may be heavier (I mean, it may have bigger mass per unit thrust), thus reducing overall plane efficiency.
- 277 *The longest recorded non-stop flight by a bird...*
New Scientist 2492. "Bar-tailed godwit is king of the skies." 26 March, 2005.
 11 September, 2007: Godwit flies 11 500 km non-stop from Alaska to New Zealand. [2qbquv]
- 278 *Optimizing hop lengths: the sweet spot is when the hops are about 5000 km long*. Source: Green (2006).
- 280 *Data for a passenger-carrying catamaran*. From [5h6xph]: Displacement (full load) 26.3 tons. On a 1050 nautical mile voyage she consumed just 4780 litres of fuel. I reckon that's a weight-transport-cost of 0.93 kWh per ton-km. I'm counting the total weight of the vessel here, by the way. The same vessel's *passenger*-transport-efficiency is roughly 35 kWh per 100 p-km.
- 281 *The Lun ekranoplan*. Sources: www.fas.org [4p3yco], (Taylor, 2002a).

Further reading: Tennekes (1997), Shyy et al. (1999).

D Solar II

On p42 we listed four solar biomass options:

1. “Coal substitution.”
2. “Petroleum substitution.”
3. Food for humans or other animals.
4. Incineration of agricultural by-products.

We’ll estimate the maximum plausible contribution of each of these processes in turn. In practice, many of these methods require so much energy to be put *in* along the way that they are scarcely net contributors (figure 6.14). But in what follows, I’ll ignore such embodied-energy costs.

Energy crops as a coal substitute

If we grow in Britain energy crops such as willow, miscanthus, or poplar (which have an average power of 0.5 W per square metre of land), then shove them in a 40%-efficient power station, the resulting power per unit area is 0.2 W/m^2 . If one eighth of Britain (500 m^2 per person) were covered in these plantations, the resulting power would be $2.5\text{ kWh/d per person}$.

Petroleum substitution

There are several ways to turn plants into liquid fuels. I’ll express the potential of each method in terms of its power per unit area (as in figure 6.11).

Britain’s main biodiesel crop, rape

Typically, rape is sown in September and harvested the following August. Currently 450 000 hectares of oilseed rape are grown in the UK each year. (That’s 2% of the UK.) Fields of rape produce 1200 litres of biodiesel per hectare per year; biodiesel has an energy of 9.8 kWh per litre; So that’s a power per unit area of 0.13 W/m^2 .

If we used 25% of Britain for oilseed rape, we’d obtain biodiesel with an energy content of $3.1\text{ kWh/d per person}$.

Sugar beet to ethanol

Sugar beet, in the UK, delivers an impressive yield of 53 t per hectare per year. And 1 t of sugar beet makes 108 litres of bioethanol. Bioethanol has an energy density of 6 kWh per litre, so this process has a power per unit area of 0.4 W/m^2 , not accounting for energy inputs required.



Figure D.1. Two trees.



Figure D.2. Oilseed rape. If used to create biodiesel, the power per unit area of rape is 0.13 W/m^2 . Photo by Tim Dunne.

Bioethanol from sugar cane

Where sugar cane can be produced (e.g., Brazil) production is 80 tons per hectare per year, which yields about 17 600 l of ethanol. Bioethanol has an energy density of 6 kWh per litre, so this process has a power per unit area of 1.2 W/m².

Bioethanol from corn in the USA

The power per unit area of bioethanol from corn is astonishingly low. Just for fun, let's report the numbers first in archaic units. 1 acre produces 122 bushels of corn per year, which makes 122 × 2.6 US gallons of ethanol, which at 84 000 BTU per gallon means a power per unit area of just 0.02 W/m² – and we haven't taken into account any of the energy losses in processing!

Cellulosic ethanol from switchgrass

Cellulosic ethanol – the wonderful “next generation” biofuel? Schmer et al. (2008) found that the net energy yield of switchgrass grown over five years on marginal cropland on 10 farms in the midcontinental US was 60 GJ per hectare per year, which is 0.2 W/m². “This is a baseline study that represents the genetic material and agronomic technology available for switchgrass production in 2000 and 2001, when the fields were planted. Improved genetics and agronomics may further enhance energy sustainability and biofuel yield of switchgrass.”

Jatropha also has low power per unit area

Jatropha is an oil-bearing crop that grows best in dry tropical regions (300–1000 mm rain per year). It likes temperatures 20–28 °C. The projected yield in hot countries on good land is 1600 litres of biodiesel per hectare per year. That's a power per unit area of 0.18 W/m². On wasteland, the yield is 583 litres per hectare per year. That's 0.065 W/m².

If people decided to use 10% of Africa to generate 0.065 W/m², and shared this power between six billion people, what would we all get? 0.8 kWh/d/p. For comparison, world oil consumption is 80 million barrels per day, which, shared between six billion people, is 23 kWh/d/p. So even if *all* of Africa were covered with jatropha plantations, the power produced would be only one third of world oil consumption.

What about algae?

Algae are just plants, so everything I've said so far applies to algae. Slimy underwater plants are no more efficient at photosynthesis than their terrestrial cousins. But there is one trick that I haven't discussed, which is

	energy density (kWh/kg)
softwood	
– air dried	4.4
– oven dried	5.5
hardwood	
– air dried	3.75
– oven dried	5.0
white office paper	4.0
glossy paper	4.1
newspaper	4.9
cardboard	4.5
coal	8
straw	4.2
poultry litter	2.4
general indust'l waste	4.4
hospital waste	3.9
municipal solid waste	2.6
refuse-derived waste	5.1
tyres	8.9

Table D.3. Calorific value of wood and similar things. Sources: Yaros (1997); Ucuncu (1993), Digest of UK Energy Statistics 2005.



standard practice in the algae-to-biodiesel community: they grow their algae in water heavily enriched with carbon dioxide, which might be collected from power stations or other industrial facilities. It takes much less effort for plants to photosynthesize if the carbon dioxide has already been concentrated for them. In a sunny spot in America, in ponds fed with concentrated CO₂ (concentrated to 10%), Ron Putt of Auburn University says that algae can grow at 30 g per square metre per day, producing 0.01 litres of biodiesel per square metre per day. This corresponds to a power per unit pond area of 4 W/m^2 – similar to the Bavaria photovoltaic farm. If you wanted to drive a typical car (doing 12 km per litre) a distance of 50 km per day, then you'd need 420 square metres of algae-ponds just to power your car; for comparison, the area of the UK per person is 4000 square metres, of which 69 m² is water (figure 6.8). Please don't forget that it's essential to feed these ponds with concentrated carbon dioxide. So this technology would be limited both by land area – how much of the UK we could turn into algal ponds – and by the availability of concentrated CO₂, the capture of which would have an energy cost (a topic discussed in Chapters 23 and 31). Let's check the limit imposed by the concentrated CO₂. To grow 30 g of algae per m² per day would require at least 60 g of CO₂ per m² per day (because the CO₂ molecule has more mass per carbon atom than the molecules in algae). If all the CO₂ from all UK power stations were captured (roughly 2½ tons per year per person), it could service 230 square metres per person of the algal ponds described above – roughly 6% of the country. This area would deliver biodiesel with a power of 24 kWh per day per person, assuming that the numbers for sunny America apply here. A plausible vision? Perhaps on one tenth of that scale? I'll leave it to you to decide.

What about algae in the sea?

Remember what I just said: the algae-to-biodiesel posse always feed their algae concentrated CO₂. If you're going out to sea, presumably pumping CO₂ into it won't be an option. And without the concentrated CO₂, the productivity of algae drops 100-fold. For algae in the sea to make a difference, a country-sized harvesting area in the sea would be required.

What about algae that produce hydrogen?

Trying to get slime to produce hydrogen in sunlight is a smart idea because it cuts out a load of chemical steps normally performed by carbohydrate-producing plants. Every chemical step reduces efficiency a little. Hydrogen can be produced directly by the photosynthetic system, right at step one. A research study from the National Renewable Energy Laboratory in Colorado predicted that a reactor filled with genetically-modified green algae, covering an area of 11 hectares in the Arizona desert, could

produce 300 kg of hydrogen per day. Hydrogen contains 39 kWh per kg, so this algae-to-hydrogen facility would deliver a power per unit area of 4.4 W/m^2 . Taking into account the estimated electricity required to run the facility, the net power delivered would be reduced to 3.6 W/m^2 . That strikes me as still quite a promising number – compare it with the Bavarian solar photovoltaic farm, for example (5 W/m^2).

Food for humans or other animals

Grain crops such as wheat, oats, barley, and corn have an energy density of about 4 kWh per kg. In the UK, wheat yields of 7.7 tons per hectare per year are typical. If the wheat is eaten by an animal, the power per unit area of this process is 0.34 W/m^2 . If 2800 m^2 of Britain (that's all agricultural land) were devoted to the growth of crops like these, the chemical energy generated would be about $24 \text{ kWh/d per person}$.

Incineration of agricultural by-products

We found a moment ago that the power per unit area of a biomass power station burning the best energy crops is 0.2 W/m^2 . If instead we grow crops for food, and put the left-overs that we don't eat into a power station – or if we feed the food to chickens and put the left-overs that come out of the chickens' back ends into a power station – what power could be delivered per unit area of farmland? Let's make a rough guess, then take a look at some real data. For a wild guess, let's imagine that by-products are harvested from half of the area of Britain (2000 m^2 per person) and trucked to power stations, and that general agricultural by-products deliver 10% as much power per unit area as the best energy crops: 0.02 W/m^2 . Multiplying this by 2000 m^2 we get $1 \text{ kWh per day per person}$.

Have I been unfair to agricultural garbage in making this wild guess? We can re-estimate the plausible production from agricultural left-overs by scaling up the prototype straw-burning power station at Elean in East Anglia. Elean's power output is 36 MW, and it uses 200 000 tons per year from land located within a 50-mile radius. If we assume this density can be replicated across the whole country, the Elean model offers 0.002 W/m^2 . At 4000 m^2 per person, that's 8 W per person, or $0.2 \text{ kWh/day per person}$.

Let's calculate this another way. UK straw production is 10 million tons per year, or 0.46 kg per day per person. At 4.2 kWh per kg, this straw has a chemical energy of 2 kWh per day per person. If all the straw were burned in 30%-efficient power stations – a proposal that wouldn't go down well with farm animals, who have other uses for straw – the electricity generated would be $0.6 \text{ kWh/d per person}$.

Landfill methane gas

At present, much of the methane gas leaking out of rubbish tips comes from biological materials, especially waste food. So, as long as we keep throwing away things like food and newspapers, landfill gas is a sustainable energy source – plus, burning that methane might be a good idea from a climate-change perspective, since methane is a stronger greenhouse-gas than CO₂. A landfill site receiving 7.5 million tons of household waste per year can generate 50 000 m³ per hour of methane.

In 1994, landfill methane emissions were estimated to be 0.05 m³ per person per day, which has a chemical energy of 0.5 kWh/d per person, and would generate 0.2 kWh(e)/d per person, if it were all converted to electricity with 40% efficiency. Landfill gas emissions are declining because of changes in legislation, and are now roughly 50% lower.

Burning household waste

SELCHP (“South East London Combined Heat and Power”) [www.selchp.com] is a 35 MW power station that is paid to burn 420 kt per year of black-bag waste from the London area. They burn the waste as a whole, without sorting. After burning, ferrous metals are removed for recycling, hazardous wastes are filtered out and sent to a special landfill site, and the remaining ash is sent for reprocessing into recycled material for road building or construction use. The calorific value of the waste is 2.5 kWh/kg, and the thermal efficiency of the power station is about 21%, so each 1 kg of waste gets turned into 0.5 kWh of electricity. The carbon emissions are about 1000 g CO₂ per kWh. Of the 35 MW generated, about 4 MW is used by the plant itself to run its machinery and filtering processes.

Scaling this idea up, if every borough had one of these, and if everyone sent 1 kg per day of waste, then we’d get 0.5 kWh(e) per day per person from waste incineration.

This is similar to the figure estimated above for methane capture at landfill sites. And remember, we can’t have both. More waste incineration means less methane gas leaking out of landfill sites. See figure 27.2, p206, and figure 27.3, p207, for further data on waste incineration.

Notes and further reading

page no.

- 283 *The power per unit area of using willow, miscanthus, or poplar, for electricity is 0.2 W/m².* Source: Select Committee on Science and Technology Minutes of Evidence – Memorandum from the Biotechnology & Biological Sciences Research Council [www.publications.parliament.uk/pa/ld200304/ldselect/ldsctech/126/4032413.htm]. “Typically a sustainable crop of 10





Figure D.4. SELCHP – your trash is their business.

dry t/ha/y of woody biomass can be produced in Northern Europe. ... Thus an area of 1 km² will produce 1000 dry t/y – enough for a power output 150 kWe at low conversion efficiencies or 300 kWe at high conversion efficiencies.” This means 0.15–0.3 W(e)/m². See also Layzell et al. (2006), [3ap71c].

283 *Oilseed rape*. Sources: Bayer Crop Science (2003), Evans (2007), www.defra.gov.uk.

– *Sugar beet*. Source: statistics.defra.gov.uk/esg/default.asp

284 *Bioethanol from corn*. Source: Shapouri et al. (1995).

– *Bioethanol from cellulose*. See also Mabee et al. (2006).

– *Jatropha*. Sources: Francis et al. (2005), Asselbergs et al. (2006).

285 *In America, in ponds fed with concentrated CO₂, algae can grow at 30 grams per square metre per day, producing 0.01 litres of biodiesel per square metre per day*. Source: Putt (2007). This calculation has ignored the energy cost of running the algae ponds and processing the algae into biodiesel. Putt describes the energy balance of a proposed design for a 100-acre algae farm, powered by methane from an animal litter digester. The farm described would in fact produce less power than the methane power input. The 100-acre farm would use 2600 kW of methane, which corresponds to an input power density of 6.4 W/m². To recap, the power density of the output, in the form of biodiesel, would be just 4.2 W/m². All proposals to make biofuels should be approached with a critical eye!

286 *A research study from the National Renewable Energy Laboratory predicted that genetically-modified green algae, covering an area of 11 hectares, could produce 300 kg of hydrogen per day*. Source: Amos (2004).

– *Elean power station*. Source: Government White Paper (2003). Elean Power Station (36 MW) – the UK’s first straw-fired power plant. *Straw production*: www.biomassenergycentre.org.uk.

287 *Landfill gas*. Sources: Matthew Chester, City University, London, personal communication; Meadows (1996), Aitchison (1996); Alan Rosevear, UK Representative on Methane to Markets Landfill Gas Sub-Committee, May 2005 [4hamks].



E Heating II

A perfectly sealed and insulated building would hold heat for ever and thus would need no heating. The two dominant reasons why buildings lose heat are:

- 1. **Conduction** – heat flowing directly through walls, windows and doors;
- 2. **Ventilation** – hot air trickling out through cracks, gaps, or deliberate ventilation ducts.

In the standard model for heat loss, both these heat flows are proportional to the temperature difference between the air inside and outside. For a typical British house, conduction is the bigger of the two losses, as we'll see.

Conduction loss

The rate of conduction of heat through a wall, ceiling, floor, or window is the product of three things: the area of the wall, a measure of conductivity of the wall known in the trade as the “U-value” or thermal transmittance, and the temperature difference –

$$\text{power loss} = \text{area} \times U \times \text{temperature difference.}$$

The U-value is usually measured in $\text{W/m}^2/\text{K}$. (One kelvin (1K) is the same as one degree Celsius (1°C).) Bigger U-values mean bigger losses of power. The thicker a wall is, the smaller its U-value. Double-glazing is about as good as a solid brick wall. (See table E.2.)

The U-values of objects that are “in series,” such as a wall and its inner lining, can be combined in the same way that electrical conductances combine:

$$u_{\text{series combination}} = 1 \bigg/ \left(\frac{1}{u_1} + \frac{1}{u_2} \right).$$

There's a worked example using this rule on page 296.

Ventilation loss

To work out the heat required to warm up incoming cold air, we need the heat capacity of air: $1.2 \text{ kJ/m}^3/\text{K}$.

In the building trade, it's conventional to describe the power-losses caused by ventilation of a space as the product of the number of changes N of the air per hour, the volume V of the space in cubic metres, the heat capacity C , and the temperature difference ΔT between the inside and



kitchen	2
bathroom	2
lounge	1
bedroom	0.5

Table E.1. Air changes per hour: typical values of N for draught-proofed rooms. The worst draughty rooms might have $N = 3$ air changes per hour. The recommended minimum rate of air exchange is between 0.5 and 1.0 air changes per hour, providing adequate fresh air for human health, for safe combustion of fuels and to prevent damage to the building fabric from excess moisture in the air (EST 2003).

	U-values (W/m ² /K)		
	old buildings	modern standards	best methods
Walls		0.45–0.6	0.12
solid masonry wall	2.4		
outer wall: 9 inch solid brick	2.2		
11 in brick-block cavity wall, unfilled	1.0		
11 in brick-block cavity wall, insulated	0.6		
Floors		0.45	0.14
suspended timber floor	0.7		
solid concrete floor	0.8		
Roofs		0.25	0.12
flat roof with 25 mm insulation	0.9		
pitched roof with 100mm insulation	0.3		
Windows			1.5
single-glazed	5.0		
double-glazed	2.9		
double-glazed, 20 mm gap	1.7		
triple-glazed	0.7–0.9		

Table E.2. U-values of walls, floors, roofs, and windows.

outside of the building.

power
(watts)

=

$C \frac{N}{1\text{ h}} V(\text{m}^3) \Delta T(\text{K})$

(E.1)

=

$(1.2\text{ kJ/m}^3/\text{K}) \frac{N}{3600\text{ s}} V(\text{m}^3) \Delta T(\text{K})$

(E.2)

=

$\frac{1}{3} NV \Delta T.$

(E.3)

Energy loss and temperature demand (degree-days)

Since energy is power $\times$ time, you can write the energy lost by *conduction* through an area in a short duration as

energy loss =

area $\times$ $U \times (\Delta T \times \text{duration}),$

and the energy lost by *ventilation* as

$\frac{1}{3} NV \times (\Delta T \times \text{duration}).$

Both these energy losses have the form

Something $\times (\Delta T \times \text{duration}),$




Figure E.3. U-values required by British and Swedish building regulations.

where the “*Something*” is measured in watts per °C. As day turns to night, and seasons pass, the temperature difference ΔT changes; we can think of a long period as being chopped into lots of small durations, during each of which the temperature difference is roughly constant. From duration to duration, the temperature difference changes, but the Somethings don’t change. When predicting a space’s total energy loss due to conduction and ventilation over a long period we thus need to multiply two things:

1. the sum of all the *Somethings* (adding $\text{area} \times U$ for all walls, roofs, floors, doors, and windows, and $\frac{1}{3}NV$ for the volume); and
2. the sum of all the Temperature difference $\times$ duration factors (for all the durations).

The first factor is a property of the building measured in watts per °C. I’ll call this the *leakiness* of the building. (This leakiness is sometimes called the building’s *heat-loss coefficient*.) The second factor is a property of the weather; it’s often expressed as a number of “degree-days,” since temperature difference is measured in degrees, and days are a convenient unit for thinking about durations. For example, if your house interior is at 18 °C, and the outside temperature is 8 °C for a week, then we say that that




Figure E.4. The temperature demand in Cambridge, 2006, visualized as an area on a graph of daily average temperatures. (a) Thermostat set to 20 °C, including cooling in summer; (b) winter thermostat set to 17 °C.

week contributed $10 \times 7 = 70$ degree-days to the $(\Delta T \times \text{duration})$ sum. I'll call the sum of all the $(\Delta T \times \text{duration})$ factors the *temperature demand* of a period.

$$\text{energy lost} = \text{leakiness} \times \text{temperature demand}.$$

We can reduce our energy loss by reducing the leakiness of the building, or by reducing our temperature demand, or both. The next two sections look more closely at these two factors, using a house in Cambridge as a case-study.

There is a third factor we must also discuss. The lost energy is replenished by the building's heating system, and by other sources of energy such as the occupants, their gadgets, their cookers, and the sun. Focussing on the heating system, the energy *delivered* by the heating is not the same as the energy *consumed* by the heating. They are related by the *coefficient of performance* of the heating system.

$$\text{energy consumed} = \text{energy delivered} / \text{coefficient of performance}.$$

For a condensing boiler burning natural gas, for example, the coefficient of performance is 90%, because 10% of the energy is lost up the chimney.

To summarise, we can reduce the energy consumption of a building in three ways:

1. by reducing temperature demand;
2. by reducing leakiness; or
3. by increasing the coefficient of performance.

We now quantify the potential of these options. (A fourth option – increasing the building's incidental heat gains, especially from the sun – may also be useful, but I won't address it here.)




Figure E.5. Temperature demand in Cambridge, in degree-days per year, as a function of thermostat setting (°C). Reducing the winter thermostat from 20 °C to 17 °C reduces the temperature demand of heating by 30%, from 3188 to 2265 degree-days. Raising the summer thermostat from 20 °C to 23 °C reduces the temperature demand of cooling by 82%, from 91 to 16 degree-days.

Temperature demand

We can visualize the temperature demand nicely on a graph of external temperature versus time (figure E.4). For a building held at a temperature of 20°C , the total temperature demand is the *area* between the horizontal line at 20°C and the external temperature. In figure E.4a, we see that, for one year in Cambridge, holding the temperature at 20°C year-round had a temperature demand of 3188 degree-days of heating and 91 degree-days of cooling. These pictures allow us easily to assess the effect of turning down the thermostat and living without air-conditioning. Turning the winter thermostat down to 17°C , the temperature demand for heating drops from 3188 degree-days to 2265 degree-days (figure E.4b), which corresponds to a 30% reduction in heating demand. Turning the thermostat down to 15°C reduces the temperature demand from 3188 to 1748 degree days, a 45% reduction.

These calculations give us a ballpark indication of the benefit of turning down thermostats, but will give an exact prediction only if we take into account two details: first, buildings naturally absorb energy from the sun, boosting the inside above the outside temperature, even without any heating; and second, the occupants and their gadget companions emit heat, so further cutting down the artificial heating requirements. The temperature demand of a location, as conventionally expressed in degree-days, is a bit of an unwieldy thing. I find it hard to remember numbers like “3500 degree-days.” And academics may find the degree-day a distressing unit, since they already have another meaning for degree days (one involving dressing up in gowns and mortar boards). We can make this quantity more meaningful and perhaps easier to work with by dividing it by 365, the number of days in the year, obtaining the temperature demand in “degree-days per day,” or, if you prefer, in plain “degrees.” Figure E.6 shows this replotted temperature demand. Expressed this way, the temperature demand is simply the *average* temperature difference between inside and outside. The highlighted temperature demands are: 8.7°C , for a thermostat setting of 20°C ; 6.2°C , for a setting of 17°C ; and 4.8°C , for a setting of 15°C .

Leakiness – example: my house

My house is a three-bedroom semi-detached house built about 1940 (figure E.7). By 2006, its kitchen had been slightly extended, and most of the windows were double-glazed. The front door and back door were both still single-glazed.

My estimate of the leakiness in 2006 is built up as shown in table E.8. The total leakiness of the house was $322\text{ W}/^{\circ}\text{C}$ (or $7.7\text{ kWh}/\text{d}/^{\circ}\text{C}$), with conductive leakiness accounting for 72% and ventilation leakiness for 28% of the total. The conductive leakiness is roughly equally divided into three parts: windows; walls; and floor and ceiling.




Figure E.6. The temperature demand in Cambridge, 2006, replotted in units of degree-days per day, also known as degrees. In these units, the temperature demand is just the average of the temperature difference between inside and outside.



Figure E.7. My house.

CONDUCTIVE LEAKINESS		area (m ²)	U-value (W/m ² /°C)	leakiness (W/°C)
Horizontal surfaces				
Pitched roof		48	0.6	28.8
Flat roof		1.6	3	4.8
Floor		50	0.8	40
Vertical surfaces				
Extension walls		24.1	0.6	14.5
Main walls		50	1	50
Thin wall (5in)		2	3	6
Single-glazed doors and windows		7.35	5	36.7
Double-glazed windows		17.8	2.9	51.6
Total conductive leakiness				232.4

VENTILATION LEAKINESS	volume (m ³)	N (air-changes per hour)	leakiness (W/°C)
Bedrooms	80	0.5	13.3
Kitchen	36	2	24
Hall	27	3	27
Other rooms	77	1	25.7
Total ventilation leakiness			90

Table E.8. Breakdown of my house’s conductive leakiness, and its ventilation leakiness, pre-2006. I’ve treated the central wall of the semi-detached house as a perfect insulating wall, but this may be wrong if the gap between the adjacent houses is actually well-ventilated. I’ve highlighted the parameters that I altered after 2006, in modifications to be described shortly.

To compare the leakinesses of two buildings that have different floor areas, we can divide the leakiness by the floor area; this gives the *heat-loss parameter* of the building, which is measured in W/°C/m². The heat-loss parameter of this house (total floor area 88 m²) is

$$3.7 \text{ W/°C/m}^2.$$

Let’s use these figures to estimate the house’s daily energy consumption on a cold winter’s day, and year-round.

On a cold day, assuming an external temperature of −1 °C and an internal temperature of 19 °C, the temperature difference is $\Delta T = 20 \text{ °C}$. If this difference is maintained for 6 hours per day then the energy lost per day is

$$322 \text{ W/°C} \times 120 \text{ degree-hours} \simeq 39 \text{ kWh.}$$

If the temperature is maintained at 19 °C for 24 hours per day, the energy lost per day is

$$155 \text{ kWh/d.}$$

To get a year-round heat-loss figure, we can take the temperature demand of Cambridge from figure E.5. With the thermostat at 19 °C, the

temperature demand in 2006 was 2866 degree-days. The average rate of heat loss, if the house is always held at 19 °C, is therefore:

$$7.7\text{ kWh/d/}^\circ\text{C} \times 2866\text{ degree-days/y}/(365\text{ days/y}) = 61\text{ kWh/d.}$$

Turning the thermostat down to 17 °C, the average rate of heat loss drops to 48 kWh/d. Turning it up to a tropical 21 °C, the average rate of heat loss is 75 kWh/d.

Effects of extra insulation

During 2007, I made the following modifications to the house:

1. Added cavity-wall insulation (which was missing in the main walls of the house) – figure 21.5.
2. Increased the insulation in the roof.
3. Added a new front door outside the old – figure 21.6.
4. Replaced the back door with a double-glazed one.
5. Double-glazed the one window that was still single-glazed.

What’s the predicted change in heat loss?

The total leakiness before the changes was 322 W/°C.

Adding cavity-wall insulation (new U-value 0.6) to the main walls reduces the house’s leakiness by 20 W/°C. The improved loft insulation (new U-value 0.3) should reduce the leakiness by 14 W/°C. The glazing modifications (new U-value 1.6–1.8) should reduce the conductive leakiness by 23 W/°C, and the ventilation leakiness by something like 24 W/°C. That’s a total reduction in leakiness of 25%, from roughly 320 to 240 W/°C (7.7 to 6 kWh/d/°C). Table E.9 shows the predicted savings from each of the modifications.

The heat-loss parameter of this house (total floor area 88m²) is thus hopefully reduced by about 25%, from 3.7 to 2.7 W/°C/m². (This is a long way from the 1.1 W/°C/m² required of a “sustainable” house in the new building codes.)

– Cavity-wall insulation (applicable to two-thirds of the wall area)	4.8 kWh/d
– Improved roof insulation	3.5 kWh/d
– Reduction in conduction from double-glazing two doors and one window	1.9 kWh/d
– Ventilation reductions in hall and kitchen from improvements to doors and windows	2.9 kWh/d

Table E.9. Break-down of the predicted reductions in heat loss from my house, on a cold winter day.

It's frustratingly hard to make a really big dent in the leakiness of an already-built house! As we saw a moment ago, a much easier way of achieving a big dent in heat loss is to turn the thermostat down. Turning down from 20 to 17°C gave a reduction in heat loss of 30%.

Combining these two actions – the physical modifications and the turning-down of the thermostat – this model predicts that heat loss should be reduced by nearly 50%. Since some heat is generated in a house by sunshine, gadgets, and humans, the reduction in gas consumption should be more than 50%.

I made all these changes to my house and monitored my meters every week. I can confirm that my heating bill indeed went down by more than 50%. As figure 21.4 showed, my gas consumption has gone down from 40 kWh/d to 13 kWh/d – a reduction of 67%.

Leakiness reduction by internal wall-coverings

Can you reduce your walls' leakiness by covering the *inside* of the wall with insulation? The answer is yes, but there may be two complications. First, the thickness of internal covering is bigger than you might expect. To transform an existing nine-inch solid brick wall (U-value 2.2 W/m²/K) into a decent 0.30 W/m²/K wall, roughly 6 cm of insulated lining board is required. [65h3cb] Second, condensation may form on the hidden surface of such internal insulation layers, leading to damp problems.

If you're not looking for such a big reduction in wall leakiness, you can get by with a thinner internal covering. For example, you can buy 1.8-cm-thick insulated wallboards with a U-value of 1.7 W/m²/K. With these over the existing wall, the U-value would be reduced from 2.2 W/m²/K to:

$$1 / \left(\frac{1}{2.2} + \frac{1}{1.7} \right) \simeq 1 \text{ W/m}^2/\text{K}.$$

Definitely a worthwhile reduction.

Air-exchange

Once a building is really well insulated, the principal loss of heat will be through ventilation (air changes) rather than through conduction. The heat loss through ventilation can be reduced by transferring the heat from the outgoing air to the incoming air. Remarkably, a great deal of this heat can indeed be transferred without any additional energy being required. The trick is to use a nose, as discovered by natural selection. A nose warms incoming air by cooling down outgoing air. There's a temperature gradient along the nose; the walls of a nose are coldest near the nostrils. The longer your nose, the better it works as a counter-current heat exchanger. In nature's noses, the direction of the air-flow usually alternates. Another way to organize a nose is to have two air-passages, one for in-flow and

one for out-flow, separate from the point of view of air, but tightly coupled with each other so that heat can easily flow between the two passages. This is how the noses work in buildings. It's conventional to call these noses heat-exchangers.

An energy-efficient house

In 1984, an energy consultant, Alan Foster, built an energy-efficient house near Cambridge; he kindly gave me his thorough measurements. The house is a timber-framed bungalow based on a Scandinavian “Heatkeeper Serrekunda” design (figure E.10), with a floor area of 140 m^2 , composed of three bedrooms, a study, two bathrooms, a living room, a kitchen, and a lobby. The wooden outside walls were supplied in kit form by a Scottish company, and the main parts of the house took only a few days to build.

The walls are 30 cm thick and have a U-value of $0.28\text{ W/m}^2/\text{°C}$. From the inside out, they consist of 13 mm of plasterboard, 27 mm airspace, a vapour barrier, 8 mm of plywood, 90 mm of rockwool, 12 mm of bitumen-impregnated fibreboard, 50 mm cavity, and 103 mm of brick. The ceiling construction is similar with 100–200 mm of rockwool insulation. The ceiling has a U-value of $0.27\text{ W/m}^2/\text{°C}$, and the floor, $0.22\text{ W/m}^2/\text{°C}$. The windows are double-glazed (U-value $2\text{ W/m}^2/\text{°C}$), with the inner panes' outer surfaces specially coated to reduce radiation. The windows are arranged to give substantial solar gain, contributing about 30% of the house's space-heating.

The house is well sealed, every door and window lined with neoprene gaskets. The house is heated by warm air pumped through floor grilles; in winter, pumps remove used air from several rooms, exhausting it to the outside, and they take in air from the loft space. The incoming air and outgoing air pass through a heat exchanger (figure E.11), which saves 60% of the heat in the extracted air. The heat exchanger is a passive device, using no energy: it's like a big metal nose, warming the incoming air with the outgoing air. On a cold winter's day, the outside air temperature was -8 °C , the temperature in the loft's air intake was 0 °C , and the air coming out of the heat exchanger was at $+8\text{ °C}$.

For the first decade, the heat was supplied entirely by electric heaters, heating a 150-gallon heat store during the overnight economy period. More recently a gas supply was brought to the house, and the space heating is now obtained from a condensing boiler.

The heat loss through conduction and ventilation is 4.2 kWh/d/°C . The *heat loss parameter* (the leakiness per square metre of floor area) is $1.25\text{ W/m}^2/\text{°C}$ (cf. my house's 2.7 W/°C/m^2).

With the house occupied by two people, the average space-heating consumption, with the thermostat set at 19 or 20 °C during the day, was 8100 kWh per year, or 22 kWh/d ; the total energy consumption for all purposes was about $15\,000\text{ kWh}$ per year, or 40 kWh/d . Expressed as an aver-




Figure E.10. The Heatkeeper Serrekunda.



Figure E.11. The Heatkeeper's heat-exchanger.

age power per unit area, that's 6.6 W/m^2 .

Figure E.12 compares the power consumption per unit area of this Heatkeeper house with my house (before and after my efficiency push) and with the European average. My house's post-efficiency-push consumption is close to that of the Heatkeeper, thanks to the adoption of lower thermostat settings.

Benchmarks for houses and offices

The German Passivhaus standard aims for power consumption for heating and cooling of $15 \text{ kWh/m}^2/\text{y}$, which is 1.7 W/m^2 ; and total power consumption of $120 \text{ kWh/m}^2/\text{y}$, which is 13.7 W/m^2 .

The average energy consumption of the UK service sector, per unit floor area, is 30 W/m^2 .

An energy-efficient office


The National Energy Foundation built themselves a low-cost low-energy building. It has solar panels for hot water, solar photovoltaic (PV) panels generating up to 6.5 kW of electricity, and is heated by a 14-kW ground-source heat pump and occasionally by a wood stove. The floor area is 400 m^2 and the number of occupants is about 30. It is a single-storey building. The walls contain 300 mm of rockwool insulation. The heat pump's coefficient of performance in winter was 2.5. The energy used is 65 kWh per year per square metre of floor area (7.4 W/m^2). The PV system delivers almost 20% of this energy.

Contemporary offices

New office buildings are often hyped up as being amazingly environment-friendly. Let's look at some numbers.

The William Gates building at Cambridge University holds computer science researchers, administrators, and a small café. Its area is $11\,110 \text{ m}^2$, and its energy consumption is 2392 MWh/y . That's a power per unit area of $215 \text{ kWh/m}^2/\text{y}$, or 25 W/m^2 . This building won a RIBA award in 2001 for its predicted energy consumption. "The architects have incorporated many environmentally friendly features into the building." [5dhups]


But are these buildings impressive? Next door, the Rutherford building, built in the 1970s without any fancy eco-claims – indeed without even double glazing – has a floor area of 4998 m^2 and consumes 1557 MWh per year; that's 0.85 kWh/d/m^2 , or 36 W/m^2 . So the award-winning building is just 30% better, in terms of power per unit area, than its simple 1970s cousin. Figure E.12 compares these buildings and another new building, the Law Faculty, with the Old Schools, which are ancient offices built pre-



Old Schools



Rutherford building



Law faculty



Gates building




Figure E.13. Ideal heat pump efficiencies. Top left: ideal electrical energy required, according to the limits of thermodynamics, to pump heat *out* of a place at temperature T_{in} when the heat is being pumped to a place at temperature $T_{out} = 35^{\circ}\text{C}$. Right: ideal electrical energy required to pump heat *into* a place at temperature T_{in} when the heat is being pumped from a place at temperature $T_{out} = 0^{\circ}\text{C}$. Bottom row: the efficiency is conventionally expressed as a “coefficient of performance,” which is the heat pumped per unit electrical energy. In practice, I understand that well-installed ground-source heat pumps and the best air-source heat pumps usually have a coefficient of performance of 3 or 4; however, government regulations in Japan have driven the coefficient of performance as high as 6.6.

1890. For all the fanfare, the difference between the new and the old is really quite disappointing!

Notice that the building power consumptions, per unit floor area, are in just the same units (W/m^2) as the renewable powers per unit area that we discussed on pages 43, 47, and 177. Comparing these consumption and production numbers helps us realize how difficult it is to power modern buildings entirely from on-site renewables. The power per unit area of biofuels (figure 6.11, p43) is $0.5 \text{ W}/\text{m}^2$; of wind farms, $2 \text{ W}/\text{m}^2$; of solar photovoltaics, $20 \text{ W}/\text{m}^2$ (figure 6.18, p47); only solar hot-water panels come in at the right sort of power per unit area, $53 \text{ W}/\text{m}^2$ (figure 6.3, p39).

Improving the coefficient of performance

You might think that the coefficient of performance of a condensing boiler, 90%, sounds pretty hard to beat. But it can be significantly improved upon, by heat pumps. Whereas the condensing boiler takes chemical energy and turns 90% of it into useful heat, the heat pump takes some electrical energy and uses it to *move* heat from one place to another (for example, from outside a building to inside). Usually the amount of useful heat delivered is much bigger than the amount of electricity used. A coefficient of performance of 3 or 4 is normal.

Theory of heat pumps

Here are the formulae for the ideal efficiency of a heat pump, that is, the electrical energy required per unit of heat pumped. If we are pumping heat from an outside place at temperature T_1 into a place at higher temperature T_2 , both temperatures being expressed relative to absolute zero (that is, T_2 , in kelvin, is given in terms of the Celsius temperature T_{in} , by $273.15 + T_{\text{in}}$), the ideal efficiency is:

$$\text{efficiency} = \frac{T_2}{T_2 - T_1}.$$

If we are pumping heat out from a place at temperature T_2 to a warmer exterior at temperature T_1 , the ideal efficiency is:

$$\text{efficiency} = \frac{T_2}{T_1 - T_2}.$$

These theoretical limits could only be achieved by systems that pump heat infinitely slowly. Notice that the ideal efficiency is bigger, the closer the inside temperature T_2 is to the outside temperature T_1 .

While in theory ground-source heat pumps might have better performance than air-source, because the ground temperature is usually closer than the air temperature to the indoor temperature, in practice an air-source heat pump might be the best and simplest choice. In cities, there may be uncertainty about the future effectiveness of ground-source heat pumps, because the more people use them in winter, the colder the ground gets; this thermal fly-tipping problem may also show up in the summer in cities where too many buildings use ground-source (or should I say “ground-sink”?) heat pumps for air-conditioning.

Heating and the ground

Here’s an interesting calculation to do. Imagine having solar heating panels on your roof, and, whenever the water in the panels gets above 50°C , pumping the water through a large rock under your house. When a dreary grey cold month comes along, you could then use the heat in the rock to warm your house. Roughly how big a 50°C rock would you need to hold enough energy to heat a house for a whole month? Let’s assume we’re after 24 kWh per day for 30 days and that the house is at 16°C . The heat capacity of granite is $0.195 \times 4200\text{ J/kg/K} = 820\text{ J/kg/K}$. The mass of granite required is:

$$\begin{aligned} \text{mass} &= \frac{\text{energy}}{\text{heat capacity} \times \text{temperature difference}} \\ &= \frac{24 \times 30 \times 3.6\text{ MJ}}{(820\text{ J/kg/}^\circ\text{C})(50^\circ\text{C} - 16^\circ\text{C})} \\ &= 100\,000\text{ kg,} \end{aligned}$$

100 tonnes, which corresponds to a cuboid of rock of size $6\text{ m} \times 6\text{ m} \times 1\text{ m}$.

Heat capacity:	$C = 820\text{ J/kg/K}$
Conductivity:	$\kappa = 2.1\text{ W/m/K}$
Density:	$\rho = 2750\text{ kg/m}^3$
Heat capacity per unit volume:	$C_V = 2.3\text{ MJ/m}^3/\text{K}$

Table E.14. Vital statistics for granite. (I use granite as an example of a typical rock.)

Ground storage without walls

OK, we’ve established the size of a useful ground store. But is it difficult to keep the heat in? Would you need to surround your rock cuboid with lots of insulation? It turns out that the ground itself is a pretty good insulator. A spike of heat put down a hole in the ground will spread as

$$\frac{1}{\sqrt{4\pi\kappa t}} \exp\left(-\frac{x^2}{4(\kappa/(C\rho))t}\right)$$

where κ is the conductivity of the ground, C is its heat capacity, and ρ is its density. This describes a bell-shaped curve with width

$$\sqrt{2\frac{\kappa}{C\rho}t};$$

for example, after six months ($t = 1.6 \times 10^7$ s), using the figures for granite ($C = 0.82$ kJ/kg/K, $\rho = 2500$ kg/m³, $\kappa = 2.1$ W/m/K), the width is 6 m.

Using the figures for water ($C = 4.2$ kJ/kg/K, $\rho = 1000$ kg/m³, $\kappa = 0.6$ W/m/K), the width is 2 m.

So if the storage region is bigger than 20 m × 20 m × 20 m then most of the heat stored will still be there in six months time (because 20 m is significantly bigger than 6 m and 2 m).

Limits of ground-source heat pumps

The low thermal conductivity of the ground is a double-edged sword. Thanks to low conductivity, the ground holds heat well for a long time. But on the other hand, low conductivity means that it’s not easy to shove heat in and out of the ground rapidly. We now explore how the conductivity of the ground limits the use of ground-source heat pumps.

Consider a neighbourhood with quite a high population density. Can everyone use ground-source heat pumps, without using active summer replenishment (as discussed on p152)? The concern is that if we all sucked heat from the ground at the same time, we might freeze the ground solid. I’m going to address this question by two calculations. First, I’ll work out the natural flux of energy in and out of the ground in summer and winter.

	(W/m/K)
water	0.6
quartz	8
granite	2.1
earth’s crust	1.7
dry soil	0.14

Table E.15. Thermal conductivities. For more data see table E.18, p304.




Figure E.16. The temperature in Cambridge, 2006, and a cartoon, which says the temperature is the sum of an annual sinusoidal variation between 3 °C and 20 °C, and a daily sinusoidal variation with range up to 10.3 °C. The average temperature is 11.5 °C.

If the flux we want to suck out of the ground in winter is much bigger than these natural fluxes then we know that our sucking is going to significantly alter ground temperatures, and may thus not be feasible. For this calculation, I'll assume the ground just below the surface is held, by the combined influence of sun, air, cloud, and night sky, at a temperature that varies slowly up and down during the year (figure E.16).

Response to external temperature variations

Working out how the temperature inside the ground responds, and what the flux in or out is, requires some advanced mathematics, which I've cordoned off in box E.19 (p306).

The payoff from this calculation is a rather beautiful diagram (figure E.17) that shows how the temperature varies in time at each depth. This diagram shows the answer for any material in terms of the *characteristic length-scale* z_0 (equation (E.7)), which depends on the conductivity κ and heat capacity C_V of the material, and on the frequency ω of the external temperature variations. (We can choose to look at either daily and yearly variations using the same theory.) At a depth of $2z_0$, the variations in temperature are one seventh of those at the surface, and lag them by about one third of a cycle (figure E.17). At a depth of $3z_0$, the variations in temperature are one twentieth of those at the surface, and lag them by half a cycle.

For the case of daily variations and solid granite, the characteristic length-scale is $z_0 = 0.16$ m. (So 32 cm of rock is the thickness you need to ride out external daily temperature oscillations.) For yearly variations and solid granite, the characteristic length-scale is $z_0 = 3$ m.

Let's focus on annual variations and discuss a few other materials. Characteristic length-scales for various materials are in the third column of table E.18. For damp sandy soils or concrete, the characteristic length-scale z_0 is similar to that of granite – about 2.6 m. In dry or peaty soils, the length-scale z_0 is shorter – about 1.3 m. That's perhaps good news because it means you don't have to dig so deep to find ground with a stable temperature. But it's also coupled with some bad news: the natural fluxes are smaller in dry soils.

The natural flux varies during the year and has a peak value (equation (E.9)) that is smaller, the smaller the conductivity.

For the case of solid granite, the peak flux is 8 W/m^2 . For dry soils, the peak flux ranges from 0.7 W/m^2 to 2.3 W/m^2 . For damp soils, the peak flux ranges from 3 W/m^2 to 8 W/m^2 .

What does this mean? I suggest we take a flux in the middle of these numbers, 5 W/m^2 , as a useful benchmark, giving guidance about what sort of power we could expect to extract, per unit area, with a ground-source heat pump. If we suck a flux significantly smaller than 5 W/m^2 , the perturbation we introduce to the natural flows will be small. If on the




Figure E.17. Temperature (in $^{\circ}\text{C}$) versus depth and time. The depths are given in units of the characteristic depth z_0 , which for granite and annual variations is 3 m.

At “depth 2” (6 m), the temperature is always about 11 or 12°C . At “depth 1” (3 m), it wobbles between 8 and 15°C .

other hand we try to suck a flux bigger than 5 W/m^2 , we should expect that we'll be shifting the temperature of the ground significantly away from its natural value, and such fluxes may be impossible to demand.

The population density of a typical English suburb corresponds to 160 m^2 per person (rows of semi-detached houses with about 400 m^2 per house, including pavements and streets). At this density of residential area, we can deduce that a ballpark limit for heat pump power delivery is

$$5 \text{ W/m}^2 \times 160 \text{ m}^2 = 800 \text{ W} = 19 \text{ kWh/d per person.}$$

This is uncomfortably close to the sort of power we would like to deliver in winter-time: it's plausible that our peak winter-time demand for hot air and hot water, in an old house like mine, might be 40 kWh/d per person.

This calculation suggests that in a typical suburban area, *not everyone can use ground-source heat pumps*, unless they are careful to actively dump heat back into the ground during the summer.

Let's do a second calculation, working out how much power we could steadily suck from a ground loop at a depth of $h = 2 \text{ m}$. Let's assume that we'll allow ourselves to suck the temperature at the ground loop down to $\Delta T = 5^\circ\text{C}$ below the average ground temperature at the surface, and let's assume that the surface temperature is constant. We can then deduce the heat flux from the surface. Assuming a conductivity of 1.2 W/m/K

	thermal conductivity κ (W/m/K)	heat capacity C_V (MJ/m ³ /K)	length-scale z_0 (m)	flux $A\sqrt{C_V\kappa\omega}$ (W/m ²)
Air	0.02	0.0012		
Water	0.57	4.18	1.2	5.7
Solid granite	2.1	2.3	3.0	8.1
Concrete	1.28	1.94	2.6	5.8
<i>Sandy soil</i>				
dry	0.30	1.28	1.5	2.3
50% saturated	1.80	2.12	2.9	7.2
100% saturated	2.20	2.96	2.7	9.5
<i>Clay soil</i>				
dry	0.25	1.42	1.3	2.2
50% saturated	1.18	2.25	2.3	6.0
100% saturated	1.58	3.10	2.3	8.2
<i>Peat soil</i>				
dry	0.06	0.58	1.0	0.7
50% saturated	0.29	2.31	1.1	3.0
100% saturated	0.50	4.02	1.1	5.3

Table E.18. Thermal conductivity and heat capacity of various materials and soil types, and the deduced length-scale $z_0 = \sqrt{\frac{2\kappa}{C_V\omega}}$ and peak flux $A\sqrt{C_V\kappa\omega}$ associated with annual temperature variations with amplitude $A = 8.3^\circ\text{C}$. The sandy and clay soils have porosity 0.4; the peat soil has porosity 0.8.

(typical of damp clay soil),

$$\text{Flux} = \kappa \times \frac{\Delta T}{h} = 3 \text{ W/m}^2.$$

If, as above, we assume a population density corresponding to 160 m² per person, then the maximum power per person deliverable by ground-source heat pumps, if everyone in a neighbourhood has them, is 480 W, which is 12 kWh/d per person.

So again we come to the conclusion that in a typical suburban area composed of poorly insulated houses like mine, *not everyone can use ground-source heat pumps*, unless they are careful to actively dump heat back into the ground during the summer. And in cities with higher population density, ground-source heat pumps are unlikely to be viable.

I therefore suggest air-source heat pumps are the best heating choice for most people.

Thermal mass

Does increasing the thermal mass of a building help reduce its heating and cooling bills? It depends. The outdoor temperature can vary during the day by about 10 °C. A building with large thermal mass – thick stone walls, for example – will naturally ride out those variations in temperature, and, without heating or cooling, will have a temperature close to the average outdoor temperature. Such buildings, in the UK, need neither heating nor cooling for many months of the year. In contrast, a poorly-insulated building with low thermal mass might be judged too hot during the day and too cool at night, leading to greater expenditure on cooling and heating.

However, large thermal mass is not always a boon. If a room is occupied in winter for just a couple of hours a day (think of a lecture room for example), the energy cost of warming the room up to a comfortable temperature will be greater, the greater the room's thermal mass. This extra invested heat will linger for longer in a thermally massive room, but if nobody is there to enjoy it, it's wasted heat. So in the case of infrequently-used rooms it makes sense to aim for a structure with low thermal mass, and to warm that small mass rapidly when required.

Notes and further reading

page no.

304 *Table E.18*. Sources: Bonan (2002),
www.hukseflux.com/thermalScience/thermalConductivity.html

If we assume the ground is made of solid homogenous material with conductivity κ and heat capacity C_V , then the temperature at depth z below the ground and time t responds to the imposed temperature at the surface in accordance with the diffusion equation

$$\frac{\partial T(z, t)}{\partial t} = \frac{\kappa}{C_V} \frac{\partial^2 T(z, t)}{\partial z^2}. \quad (\text{E.4})$$

For a sinusoidal imposed temperature with frequency ω and amplitude A at depth $z = 0$,

$$T(0, t) = T_{\text{surface}}(t) = T_{\text{average}} + A \cos(\omega t), \quad (\text{E.5})$$

the resulting temperature at depth z and time t is a decaying and oscillating function

$$T(z, t) = T_{\text{average}} + A e^{-z/z_0} \cos(\omega t - z/z_0), \quad (\text{E.6})$$

where z_0 is the characteristic length-scale of both the decay and the oscillation,

$$z_0 = \sqrt{\frac{2\kappa}{C_V \omega}}. \quad (\text{E.7})$$

The flux of heat (the power per unit area) at depth z is

$$\kappa \frac{\partial T}{\partial z} = \kappa \frac{A}{z_0} \sqrt{2} e^{-z/z_0} \sin(\omega t - z/z_0 - \pi/4). \quad (\text{E.8})$$

For example, at the surface, the peak flux is

$$\kappa \frac{A}{z_0} \sqrt{2} = A \sqrt{C_V \kappa \omega}. \quad (\text{E.9})$$

Box E.19. Working out the natural flux caused by sinusoidal temperature variations.

F Waves II

The physics of deep-water waves

Waves contain energy in two forms: potential energy, and kinetic energy. The potential energy is the energy required to move all the water from the troughs to the crests. The kinetic energy is associated with the water moving around.

People sometimes assume that when the crest of a wave moves across an ocean at 30 miles per hour, the water in that crest must also be moving at 30 miles per hour in the same direction. But this isn't so. It's just like a Mexican wave. When the wave rushes round the stadium, the humans who are making the wave aren't themselves moving round the stadium: they just bob up and down a little. The motion of a piece of water in the ocean is similar: if you focused on a bit of seaweed floating in the water as waves go by, you'd see that the seaweed moves up and down, and also a little to and fro in the direction of travel of the wave – the exact effect could be recreated in a Mexican wave if people moved like window-cleaners, polishing a big piece of glass in a circular motion. The wave has potential energy because of the elevation of the crests above the troughs. And it has kinetic energy because of the small circular bobbing motion of the water.

Our rough calculation of the power in ocean waves will require three ingredients: an estimate of the period T of the waves (the time between crests), an estimate of the height h of the waves, and a physics formula that tells us how to work out the speed v of the wave from its period.

The wavelength λ and period of the waves (the distance and time respectively between two adjacent crests) depend on the speed of the wind that creates the waves, as shown in figure F.1. The height of the waves doesn't depend on the windspeed; rather, it depends on how long the wind has been caressing the water surface.

You can estimate the period of ocean waves by recalling the time between waves arriving on an ocean beach. Is 10 seconds reasonable? For the height of ocean waves, let's assume an amplitude of 1 m, which means 2 m from trough to crest. In waves this high, a man in a dinghy can't see beyond the nearest crest when he's in a trough; I think this height is bigger than average, but we can revisit this estimate if we decide it's important. The speed of deep-water waves is related to the time T between crests by the physics formula (see Faber (1995), p170):

$$v = \frac{gT}{2\pi},$$

where g is the acceleration of gravity (9.8 m/s^2). For example, if $T = 10$ seconds, then $v = 16\text{ m/s}$. The wavelength of such a wave – the distance between crests – is $\lambda = vT = gT^2/2\pi = 160\text{ m}$.





Figure F.1. Facts about deep-water waves. In all four figures the horizontal axis is the wave speed in m/s. From top to bottom the graphs show: wind speed (in m/s) required to make a wave with this wave speed; period (in seconds) of a wave; wavelength (in m) of a wave; and power density (in kW/m) of a wave with amplitude 1 m.



For a wave of wavelength λ and period T , if the height of each crest and depth of each trough is $h = 1$ m, the potential energy passing per unit time, per unit length, is

$$P_{\text{potential}} \simeq m^* g \bar{h} / T, \quad (\text{F.1})$$

where m^* is the mass per unit length, which is roughly $\frac{1}{2}\rho h(\lambda/2)$ (approximating the area of the shaded crest in figure F.2 by the area of a triangle), and $\bar{h}$ is the change in height of the centre-of-mass of the chunk of elevated water, which is roughly h . So

$$P_{\text{potential}} \simeq \frac{1}{2}\rho h \frac{\lambda}{2} g h / T. \quad (\text{F.2})$$

(To find the potential energy properly, we should have done an integral here; it would have given the same answer.) Now λ/T is simply the speed at which the wave travels, v , so:

$$P_{\text{potential}} \simeq \frac{1}{4}\rho g h^2 v. \quad (\text{F.3})$$

Waves have kinetic energy as well as potential energy, and, remarkably, these are exactly equal, although I don't show that calculation here; so the total power of the waves is double the power calculated from potential

Figure F.2. A wave has energy in two forms: potential energy associated with raising water out of the light-shaded troughs into the heavy-shaded crests; and kinetic energy of all the water within a few wavelengths of the surface – the speed of the water is indicated by the small arrows. The speed of the wave, travelling from left to right, is indicated by the much bigger arrow at the top.

energy.

$$P_{\text{total}} \simeq \frac{1}{2} \rho g h^2 v. \quad (\text{F.4})$$

There's only one thing wrong with this answer: it's too big, because we've neglected a strange property of dispersive waves: the energy in the wave doesn't actually travel at the same speed as the crests; it travels at a speed called the group velocity, which for deep-water waves is *half* of the speed v . You can see that the energy travels slower than the crests by chucking a pebble in a pond and watching the expanding waves carefully. What this means is that equation (F.4) is wrong: we need to halve it. The correct power per unit length of wave-front is

$$P_{\text{total}} = \frac{1}{4} \rho g h^2 v. \quad (\text{F.5})$$

Plugging in $v = 16 \text{ m/s}$ and $h = 1 \text{ m}$, we find

$$P_{\text{total}} = \frac{1}{4} \rho g h^2 v = 40 \text{ kW/m}. \quad (\text{F.6})$$

This rough estimate agrees with real measurements in the Atlantic (Mollison, 1986). (See p75.)

The losses from viscosity are minimal: a wave of 9 seconds period would have to go three times round the world to lose 10% of its amplitude.

Real wave power systems

Deep-water devices

How effective are real systems at extracting power from waves? Stephen Salter's "duck" has been well characterized: a row of 16-m diameter ducks, feeding off Atlantic waves with an average power of 45 kW/m, would deliver 19 kW/m, including transmission to central Scotland (Mollison, 1986).

The Pelamis device, created by Ocean Power Delivery, has taken over the Salter duck's mantle as the leading floating deep-water wave device. Each snake-like device is 130 m long and is made of a chain of four segments, each 3.5 m in diameter. It has a maximum power output of 750 kW. The Pelamises are designed to be moored in a depth of about 50 m. In a wavefarm, 39 devices in three rows would face the principal wave direction, occupying an area of ocean, about 400 m long and 2.5 km wide (an area of 1 km²). The effective cross-section of a single Pelamis is 7 m (i.e., for good waves, it extracts 100% of the energy that would cross 7 m). The company says that such a wave-farm would deliver about 10 kW/m.

Shallow-water devices

Typically 70% of energy in ocean waves is lost through bottom-friction as the depth decreases from 100m to 15m. So the average wave-power per unit length of coastline in shallow waters is reduced to about 12kW/m. The Oyster, developed by Queen's University Belfast and Aquamarine Power Ltd [www.aquamarinepower.com], is a bottom-mounted flap, about 12m high, that is intended to be deployed in waters about 12m deep, in areas where the average incident wave power is greater than 15kW/m. Its peak power is 600kW. A single device would produce about 270kW in wave heights greater than 3.5m. It's predicted that an Oyster would have a bigger power per unit mass of hardware than a Pelamis.

Oysters could also be used to directly drive reverse-osmosis desalination facilities. "The peak freshwater output of an Oyster desalinator is between 2000 and 6000m³/day." That production has a value, going by the Jersey facility (which uses 8kWh per m³), equivalent to 600–2000kW of electricity.

G Tide II

Power density of tidal pools

To estimate the power of an artificial tide-pool, imagine that it's filled rapidly at high tide, and emptied rapidly at low tide. Power is generated in both directions, on the ebb and on the flood. (This is called two-way generation or double-effect generation.) The change in potential energy of the water, each six hours, is mgh , where h is the change in height of the centre of mass of the water, which is half the range. (The range is the difference in height between low and high tide; figure G.1.) The mass per unit area covered by tide-pool is $\rho \times (2h)$, where ρ is the density of water (1000 kg/m^3). So the power per unit area generated by a tide-pool is

$$\frac{2\rho gh}{6 \text{ hours}},$$

assuming perfectly efficient generators. Plugging in $h = 2 \text{ m}$ (i.e., range 4 m), we find the power per unit area of tide-pool is 3.6 W/m^2 . Allowing for an efficiency of 90% for conversion of this power to electricity, we get

$$\text{power per unit area of tide-pool} \simeq 3 \text{ W/m}^2.$$

So to generate 1 GW of power (on average), we need a tide-pool with an area of about 300 km^2 . A circular pool with diameter 20 km would do the trick. (For comparison, the area of the Severn estuary behind the proposed barrage is about 550 km^2 , and the area of the Wash is more than 400 km^2 .)

If a tide-pool produces electricity in one direction only, the power per unit area is halved. The average power density of the tidal barrage at La Rance, where the mean tidal range is 10.9 m, has been 2.7 W/m^2 for decades (p87).

The raw tidal resource

The tides around Britain are genuine tidal waves. (Tsunamis, which are called "tidal waves," have nothing to do with tides: they are caused by underwater landslides and earthquakes.) The location of the high tide (the crest of the tidal wave) moves much faster than the tidal flow – 100 miles per hour, say, while the water itself moves at just 1 mile per hour.

The energy we can extract from tides, using tidal pools or tide farms, can never be more than the energy of these tidal waves from the Atlantic. We can estimate the total power of these great Atlantic tidal waves in the same way that we estimate the power of ordinary wind-generated waves. The next section describes a standard model for the power arriving in





Figure G.1. A tide-pool in cross section. The pool was filled at high tide, and now it's low tide. We let the water out through the electricity generator to turn the water's potential energy into electricity.



travelling waves in water of depth d that is shallow compared to the wavelength of the waves (figure G.2). The power per unit length of wavecrest of shallow-water tidal waves is

$$\rho g^{3/2} \sqrt{d} h^2 / 2. \tag{G.1}$$

Table G.3 shows the power per unit length of wave crest for some plausible figures. If $d = 100$ m, and $h = 1$ or 2 m, the power per unit length of wave crest is 150 kW/m or 600 kW/m respectively. These figures are impressive compared with the raw power per unit length of ordinary Atlantic deep-water waves, 40 kW/m (Chapter F). Atlantic waves and the Atlantic tide have similar vertical amplitudes (about 1 m), but the raw power in tides is roughly 10 times bigger than that of ordinary wind-driven waves.

Taylor (1920) worked out a more detailed model of tidal power that includes important details such as the Coriolis effect (the effect produced by the earth’s daily rotation), the existence of tidal waves travelling in the opposite direction, and the direct effect of the moon on the energy flow in the Irish Sea. Since then, experimental measurements and computer models have verified and extended Taylor’s analysis. Flather (1976) built a detailed numerical model of the lunar tide, chopping the continental shelf around the British Isles into roughly 1000 square cells. Flather estimated that the total average power entering this region is 215 GW. According to his model, 180 GW enters the gap between France and Ireland. From Northern Ireland round to Shetland, the incoming power is 49 GW. Between Shetland and Norway there is a net loss of 5 GW. As shown in figure G.4, Cartwright et al. (1980) found experimentally that the average power transmission was 60 GW between Malin Head (Ireland) and Florø (Norway) and 190 GW between Valentia (Ireland) and the Brittany coast near Ouessant. The power entering the Irish Sea was found to be 45 GW, and entering the North Sea via the Dover Straits, 16.7 GW.

The power of tidal waves

This section, which can safely be skipped, provides more details behind the formula for tidal power used in the previous section. I’m going to

Figure G.2. A shallow-water wave. Just like a deep-water wave, the wave has energy in two forms: potential energy associated with raising water out of the light-shaded troughs into the heavy-shaded crests; and kinetic energy of all the water moving around as indicated by the small arrows. The speed of the wave, travelling from left to right, is indicated by the much bigger arrow at the top. For tidal waves, a typical depth might be 100 m, the crest vertical amplitude at the surface 1 or 2 m, and the water velocity amplitude 0.3 or 0.6 m/s.

h (m)	$\rho g^{3/2} \sqrt{d} h^2 / 2$ (kW/m)
0.9	125
1.0	155
1.2	220
1.5	345
1.75	470
2.0	600
2.25	780

Table G.3. Power fluxes (power per unit length of wave crest) for depth $d = 100$ m.

go into this model of tidal power in some detail because most of the official estimates of the UK tidal resource have been based on a model that I believe is incorrect.

Figure G.2 shows a model for a tidal wave travelling across relatively shallow water. This model is intended as a cartoon, for example, of tidal crests moving up the English channel or down the North Sea. It's important to distinguish the speed U at which the water itself moves (which might be about 1 mile per hour) from the speed v at which the high tide moves, which is typically 100 or 200 miles per hour.

The water has depth d . Crests and troughs of water are injected from the left hand side by the 12-hourly ocean tides. The crests and troughs move with velocity

$$v = \sqrt{gd}. \quad (\text{G.2})$$

We assume that the wavelength is much bigger than the depth, and we neglect details such as Coriolis forces and density variations in the water. Call the vertical amplitude of the tide h . For the standard assumption of nearly-vorticity-free flow, the horizontal velocity of the water is near-constant with depth. The horizontal velocity U is proportional to the surface displacement and can be found by conservation of mass:

$$U = vh/d. \quad (\text{G.3})$$

If the depth decreases, the wave velocity v reduces (equation (G.2)). For the present discussion we'll assume the depth is constant. Energy flows from left to right at some rate. How should this total tidal power be estimated? And what's the *maximum* power that could be extracted?

One suggestion is to choose a cross-section and estimate the average *flux of kinetic energy* across that plane, then assert that this quantity represents the power that could be extracted. This kinetic-energy-flux method was used by consultants Black and Veatch to estimate the UK resource. In our cartoon model, we can compute the total power by other means. We'll see that the kinetic-energy-flux answer is too small by a significant factor.

The peak kinetic-energy flux at any section is

$$K_{\text{BV}} = \frac{1}{2} \rho A U^3, \quad (\text{G.4})$$

where A is the cross-sectional area. (This is the formula for kinetic energy flux, which we encountered in Chapter B.)

The true total incident power is not equal to this kinetic-energy flux. The true total incident power in a shallow-water wave is a standard textbook calculation; one way to get it is to find the total energy present in one wavelength and divide by the period. The total energy per wavelength is the sum of the potential energy and the kinetic energy. The kinetic energy happens to be identical to the potential energy. (This is a standard feature of almost all things that wobble, be they masses on springs or children




Figure G.4. Average tidal powers measured by Cartwright et al. (1980).

on swings.) So to compute the total energy all we need to do is compute one of the two – the potential energy per wavelength, or the kinetic energy per wavelength – then double it. The potential energy of a wave (per wavelength and per unit width of wavefront) is found by integration to be

$$\frac{1}{4}\rho gh^2\lambda. \quad (\text{G.5})$$

So, doubling and dividing by the period, the true power of this model shallow-water tidal wave is

$$\text{power} = \frac{1}{2}(\rho gh^2\lambda) \times w/T = \frac{1}{2}\rho gh^2v \times w, \quad (\text{G.6})$$

where w is the width of the wavefront. Substituting $v = \sqrt{gd}$,

$$\text{power} = \rho gh^2\sqrt{gd} \times w/2 = \rho g^{3/2}\sqrt{d}h^2 \times w/2. \quad (\text{G.7})$$

Let's compare this power with the kinetic-energy flux K_{BV} . Strikingly, the two expressions scale differently with the amplitude h . Using the amplitude conversion relation (G.3), the crest velocity (G.2), and $A = wd$, we can re-express the kinetic-energy flux as

$$K_{\text{BV}} = \frac{1}{2}\rho AU^3 = \frac{1}{2}\rho wd(vh/d)^3 = \rho \left(g^{3/2}/\sqrt{d}\right) h^3 \times w/2. \quad (\text{G.8})$$

So the kinetic-energy-flux method suggests that the total power of a shallow-water wave scales as amplitude *cubed* (equation (G.8)); but the correct formula shows that the power scales as amplitude *squared* (equation (G.7)).

The ratio is

$$\frac{K_{\text{BV}}}{\text{power}} = \frac{\rho w \left(g^{3/2}/\sqrt{d}\right) h^3}{\rho g^{3/2}h^2\sqrt{d}w} = \frac{h}{d}. \quad (\text{G.9})$$

Because h is usually much smaller than d (h is about 1 m or 2 m, while d is 100 m or 10 m), estimates of tidal power resources that are based on the kinetic-energy-flux method may be *much too small*, at least in cases where this shallow-water cartoon of tidal waves is appropriate.

Moreover, estimates based on the kinetic-energy-flux method incorrectly assert that the total available power at springs (the biggest tides) is eight times greater than at neaps (the smallest tides), assuming an amplitude ratio, springs to neaps, of two to one; but the correct answer is that the total available power of a travelling wave scales as its amplitude squared, so the springs-to-neaps ratio of total-incoming-power is four to one.

Effect of shelving of sea bed, and Coriolis force

If the depth d decreases gradually and the width remains constant such that there is minimal reflection or absorption of the incoming power, then




Figure G.5. (a) Tidal current over a 21-day period at a location where the maximum current at spring tide is 2.9 knots (1.5 m/s) and the maximum current at neap tide is 1.8 knots (0.9 m/s).

(b) The power per unit sea-floor area over a nine-day period extending from spring tides to neap tides. The power peaks four times per day, and has a maximum of about 27 W/m². The average power of the tide farm is 6.4 W/m².

the power of the wave will remain constant. This means $\sqrt{dh^2}$ is a constant, so we deduce that the height of the tide scales with depth as $h \sim 1/d^{1/4}$.

This is a crude model. One neglected detail is the Coriolis effect. The Coriolis force causes tidal crests and troughs to tend to drive on the right – for example, going up the English Channel, the high tides are higher and the low tides are lower on the French side of the channel. By neglecting this effect I may have introduced some error into the estimates.

Power density of tidal stream farms

Imagine sticking underwater windmills on the sea-bed. The flow of water will turn the windmills. Because the density of water is roughly 1000 times that of air, the power of water flow is 1000 times greater than the power of wind at the same speed.

What power could tidal stream farms extract? It depends crucially on whether or not we can add up the power contributions of tidefarms on *adjacent* pieces of sea-floor. For wind, this additivity assumption is believed to work fine: as long as the wind turbines are spaced a standard distance apart from each other, the total power delivered by 10 adjacent wind farms is the sum of the powers that each would deliver if it were alone.

Does the same go for tide farms? Or do underwater windmills interfere with each other's power extraction in a different way? I don't think the answer to this question is known in general. We can name two alternative assumptions, however, and identify cartoon situations in which each assumption seems valid. The "tide is like wind" assumption says that you can put tide-turbines all over the sea-bed, spaced about 5 diameters apart from each other, and they won't interfere with each other, no matter how much of the sea-bed you cover with such tide farms.

The "you can have only one row" assumption, in contrast, asserts that the maximum power extractable in a region is the power that would be delivered by a *single* row of turbines facing the flow. A situation where this assumption is correct is the special case of a hydroelectric dam: if the water from the dam passes through a single well-designed turbine, there's no point putting any more turbines behind that one. You can't get 100

times more power by putting 99 more turbines downstream from the first. The oomph gets extracted by the first one, and there isn't any more oomph left for the others. The "you can have only one row" assumption is the right assumption for estimating the extractable power in a place where water flows through a narrow channel from approximately stationary water at one height into another body of water at a lower height. (This case is analysed by Garrett and Cummins (2005, 2007).)

I'm now going to nail my colours to a mast. I think that in many places round the British Isles, the "tide is like wind" assumption is a good approximation. Perhaps some spots have some of the character of a narrow channel. In those spots, my estimates may be over-estimates.

Let's assume that the rules for laying out a sensible tide farm will be similar to those for wind farms, and that the efficiency of the tidemills will be like that of the best windmills, about 1/2. We can then steal the formula for the power of a wind farm (per unit land area) from p265. The power per unit sea-floor area is

$$\frac{\text{power per tidemill}}{\text{area per tidemill}} = \frac{\pi}{200} \frac{1}{2} \rho U^3$$

Using this formula, table G.6 shows this tide farm power for a few tidal currents.

Now, what are typical tidal currents? Tidal charts usually give the currents associated with the tides with the largest range (called spring tides) and the tides with the smallest range (called neap tides). Spring tides occur shortly after each full moon and each new moon. Neap tides occur shortly after the first and third quarters of the moon. The power of a tide farm would vary throughout the day in a completely predictable manner. Figure G.5 illustrates the variation of power density of a tide farm with a maximum current of 1.5 m/s. The average power density of this tide farm would be 6.4 W/m². There are many places around the British Isles where the power per unit area of tide farm would be 6 W/m² or more. This power density is similar to our estimates of the power densities of wind farms (2–3 W/m²) and of photovoltaic solar farms (5–10 W/m²).

We'll now use this "tide farms are like wind farms" theory to estimate the extractable power from tidal streams in promising regions around the British Isles. As a sanity check, we'll also work out the total tidal power crossing each of these regions, using the "power of tidal waves" theory, to check our tide farm's estimated power isn't bigger than the total power available. The main locations around the British Isles where tidal currents are large are shown in figure G.7.

I estimated the typical peak currents at six locations with large currents by looking at tidal charts in *Reed's Nautical Almanac*. (These estimates could easily be off by 30%.) Have I over-estimated or under-estimated the area of each region? I haven't surveyed the sea floor so I don't know if some regions might be unsuitable in some way – too deep, or too shallow, or too

U		tide farm
(m/s)	(knots)	power (W/m ²)
0.5	1	1
1	2	8
2	4	60
3	6	200
4	8	500
5	10	1000

Table G.6. Tide farm power density (in watts per square metre of sea-floor) as a function of flow speed *U*. (1 knot = 1 nautical mile per hour = 0.514 m/s.) The power density is computed using $\frac{\pi}{200} \frac{1}{2} \rho U^3$ (equation (G.10)).




Figure G.7. Regions around the British Isles where peak tidal flows exceed 1 m/s. The six darkly-coloured regions are included in table G.8:

- 1. the English channel (south of the Isle of Wight);
- 2. the Bristol channel;
- 3. to the north of Anglesey;
- 4. to the north of the Isle of Man;
- 5. between Northern Ireland, the Mull of Kintyre, and Islay; and
- 6. the Pentland Firth (between Orkney and mainland Scotland), and within the Orkneys.

There are also enormous currents around the Channel Islands, but they are not governed by the UK. Runner-up regions include the North Sea, from the Thames (London) to the Wash (Kings Lynn). The contours show water depths greater than 100 m. Tidal data are from Reed’s Nautical Almanac and DTI Atlas of UK Marine Renewable Energy Resources (2004).

tricky to build on.

Admitting all these uncertainties, I arrive at an estimated total power of 9 kWh/d per person from tidal stream-farms. This corresponds to 9% of the raw incoming power mentioned on p83, 100 kWh per day per person. (The extraction of 1.1 kWh/d/p in the Bristol channel, region 2, might conflict with power generation by the Severn barrage; it would depend on whether the tide farm significantly adds to the existing natural friction created by the channel, or replaces it.)

Region	U		power density	area	average power			raw power	
	N	S				d	w	N	S
			(W/m ²)	(km ²)	(kWh/d/p)	(m)	(km)	(kWh/d/p)	
1	1.7	3.1	7	400	1.1	30	30	2.3	7.8
2	1.8	3.2	8	350	1.1	30	17	1.5	4.7
3	1.3	2.3	2.9	1000	1.2	50	30	3.0	9.3
4	1.7	3.4	9	400	1.4	30	20	1.5	6.3
5	1.7	3.1	7	300	0.8	40	10	1.2	4.0
6	5.0	9.0	170	50	3.5	70	10	24	78
Total					9				

(a)

(b)

Table G.8. (a) Tidal power estimates assuming that stream farms are like wind farms. The power density is the average power per unit area of sea floor. The six regions are indicated in figure G.7. N = Neaps. S = Springs. (b) For comparison, this table shows the raw incoming power estimated using equation (G.1) (p312).

v (m/s)	v (knots)	Friction power density (W/m ²)		tide farm power density (W/m ²)
		$R_1 = 0.01$	$R_1 = 0.003$	
0.5	1	1.25	0.4	1
1	2	10	3	8
2	4	80	24	60
3	6	270	80	200
4	8	640	190	500
5	10	1250	375	1000

Table G.9. Friction power density $R_1\rho U^3$ (in watts per square metre of sea-floor) as a function of flow speed, assuming $R_1 = 0.01$ or 0.003 . Flather (1976) uses $R_1 = 0.0025\text{--}0.003$; Taylor (1920) uses 0.002 . (1 knot = 1 nautical mile per hour = 0.514 m/s .) The final column shows the tide farm power estimated in table G.6. For further reading see Kowalik (2004), Sleath (1984).

Estimating the tidal resource via bottom friction

Another way to estimate the power available from tide is to compute how much power is already dissipated by friction on the sea floor. A coating of turbines placed just above the sea floor could act as a substitute bottom, exerting roughly the same drag on the passing water as the sea floor used to exert, and extracting roughly the same amount of power as friction used to dissipate, without significantly altering the tidal flows.

So, what’s the power dissipated by “bottom friction”? Unfortunately, there isn’t a straightforward model of bottom friction. It depends on the roughness of the sea bed and the material that the bed is made from – and even given this information, the correct formula to use is not settled. One widely used model says that the magnitude of the stress (force per unit area) is $R_1\rho U^2$, where U is the average flow velocity and R_1 is a dimensionless quantity called the shear friction coefficient. We can estimate the power dissipated per unit area by multiplying the stress by the velocity. Table G.9 shows the power dissipated in friction, $R_1\rho U^3$, assuming $R_1 = 0.01$ or $R_1 = 0.003$. For values of the shear friction coefficient in this range, the friction power is very similar to the estimated power that a tide farm would deliver. This is good news, because it suggests that planting a forest of underwater windmills on the sea-bottom, spaced five diameters apart, won’t radically alter the flow. The natural friction already has an effect that is in the same ballpark.

Tidal pools with pumping

“The pumping trick” artificially increases the amplitude of the tides in a tidal pool so as to amplify the power obtained. The energy cost of pumping *in* extra water at high tide is repaid with interest when the same water is let out at low tide; similarly, extra water can be pumped *out* at low tide, then let back in at high tide. The pumping trick is sometimes used at La Rance, boosting its net power generation by about 10% (Wilson and Balls, 1990). Let’s work out the theoretical limit for this technology. I’ll assume

tidal amplitude (half-range) h (m)	optimal boost height b (m)	power with pumping (W/m ²)	power without pumping (W/m ²)
1.0	6.5	3.5	0.8
2.0	13	14	3.3
3.0	20	31	7.4
4.0	26	56	13

Table G.10. Theoretical power density from tidal power using the pumping trick, assuming no constraint on the height of the basin's walls.

that generation has an efficiency of $\epsilon_g = 0.9$ and that pumping has an efficiency of $\epsilon_p = 0.85$. Let the tidal range be $2h$. I'll assume for simplicity that the prices of buying and selling electricity are the same at all times, so that the optimal height boost b to which the pool is pumped above high water is given by (marginal cost of extra pumping = marginal return of extra water):

$$b/\epsilon_p = \epsilon_g(b + 2h).$$

Defining the round-trip efficiency $\epsilon = \epsilon_g\epsilon_p$, we have

$$b = 2h \frac{\epsilon}{1 - \epsilon}.$$

For example, with a tidal range of $2h = 4$ m, and a round-trip efficiency of $\epsilon = 76\%$, the optimal boost is $b = 13$ m. This is the maximum height to which pumping can be justified if the price of electricity is constant.

Let's assume the complementary trick is used at low tide. (This requires the basin to have a vertical range of 30 m!) The delivered power per unit area is then

$$\left(\frac{1}{2} \rho g \epsilon_g (b + 2h)^2 - \frac{1}{2} \rho g \frac{1}{\epsilon_p} b^2 \right) / T,$$

where T is the time from high tide to low tide. We can express this as the maximum possible power density without pumping, $\epsilon_g 2 \rho g h^2 / T$, scaled up by a boost factor

$$\left(\frac{1}{1 - \epsilon} \right),$$

which is roughly a factor of 4. Table G.10 shows the theoretical power density that pumping could deliver. Unfortunately, this pumping trick will rarely be exploited to the full because of the economics of basin construction: full exploitation of pumping requires the total height of the pool to be roughly 4 times the tidal range, and increases the delivered power four-fold. But the amount of material in a sea-wall of height H scales as H^2 , so the cost of constructing a wall four times as high will be more than four times as big. Extra cash would probably be better spent on enlarging a tidal pool horizontally rather than vertically.

The pumping trick can nevertheless be used for free on any day when the range of natural tides is smaller than the maximum range: the water

tidal amplitude (half-range) h (m)	boost height b (m)	power with pumping (W/m ²)	power without pumping (W/m ²)
1.0	1.0	1.6	0.8
2.0	2.0	6.3	3.3
3.0	3.0	14	7.4
4.0	4.0	25	13

Table G.11. Power density offered by the pumping trick, assuming the boost height is constrained to be the same as the tidal amplitude. This assumption applies, for example, at neap tides, if the pumping pushes the tidal range up to the springs range.

level at high tide can be pumped up to the maximum. Table G.11 gives the power delivered if the boost height is set to h , that is, the range in the pool is just double the external range. A doubling of vertical range is easy at neap tides, since neap tides are typically about half as high as spring tides. Pumping the pool at neaps so that the full springs range is used thus allows neap tides to deliver roughly twice as much power as they would offer without pumping. So a system with pumping would show two-weekly variations in power of just a factor of 2 instead of 4.

Getting “always-on” tidal power by using two basins

Here’s a neat idea: have two basins, one of which is the “full” basin and one the “empty” basin; every high tide, the full basin is topped up; every low tide, the empty basin is emptied. These toppings-up and emptyings could be done either passively through sluices, or actively by pumps (using the trick mentioned above). Whenever power is required, water is allowed to flow from the full basin to the empty basin, or (better in power terms) between one of the basins and the sea. The capital cost of a two-basin scheme may be bigger because of the need for extra walls; the big win is that power is available all the time, so the facility can follow demand.

We can use power generated from the empty basin to pump extra water into the full basin at high tide, and similarly use power from the full basin to pump down the empty basin at low tide. This self-pumping would boost the total power delivered by the facility without ever needing to buy energy from the grid. It’s a delightful feature of a two-pool solution that the optimal time to *pump* water into the high pool is high tide, which is also the optimal time to *generate* power from the low pool. Similarly, low tide is the perfect time to pump down the low pool, and it’s the perfect time to generate power from the high pool. In a simple simulation, I’ve found that a two-lagoon system in a location with a natural tidal range of 4 m can, with an appropriate pumping schedule, deliver a *steady* power of 4.5 W/m² (MacKay, 2007a). One lagoon’s water level is always kept above mean sea-level; the other lagoon’s level is always kept below mean sea-level. This power density of 4.5 W/m² is 50% bigger than the maximum possible average power density of an ordinary tide-pool in the same lo-




Figure G.12. Different ways to use the tidal pumping trick. Two lagoons are located at sea-level. (a) One simple way of using two lagoons is to label one the high pool and the other the low pool; when the surrounding sea level is near to high tide, let water into the high pool, or actively pump it in (using electricity from other sources); and similarly, when the sea level is near to low tide, empty the low pool, either passively or by active pumping; then whenever power is sufficiently valuable, generate power on demand by letting water from the high pool to the low pool. (b) Another arrangement that might deliver more power per unit area has no flow of water between the two lagoons. While one lagoon is being pumped full or pumped empty, the other lagoon can deliver steady, demand-following power to the grid. Pumping may be powered by bursty sources such as wind, by spare power from the grid (say, nuclear power stations), or by the other half of the facility, using one lagoon's power to pump the other lagoon up or down.

cation (3 W/m^2). The steady power of the lagoon system would be more valuable than the intermittent and less-flexible power from the ordinary tide-pool.

A two-basin system could also function as a pumped-storage facility.

Notes

page no.

311 *Efficiency of 90%...* Turbines are about 90% efficient for heads of 3.7 m or more. Baker et al. (2006).

320 *Getting “always-on” tidal power by using two basins.* There is a two-basin tidal power plant at Haishan, Maoyan Island, China. A single generator located between the two basins, as shown in figure G.12(a), delivers power continuously, and generates 39 kW on average. [2bqapk].

Further reading: Shaw and Watson (2003b); Blunden and Bahaj (2007); Charlier (2003a,b).

For further reading on bottom friction and variation of flow with depth, see Sleath (1984).

For more on the estimation of the UK tidal resource, see MacKay (2007b).

For more on tidal lagoons, see MacKay (2007a).

H Stuff II

Imported energy

Dieter Helm and his colleagues estimated the footprint of each pound's worth of imports from country X using the average carbon intensity of country X's economy (that is, the ratio of their carbon emissions to their gross domestic product). They concluded that the embodied carbon in imports to Britain (which should arguably be added to Britain's official carbon footprint of 11 tons CO₂e per year per person) is roughly 16 tons CO₂e per year per person. A subsequent, more detailed study commissioned by DEFRA estimated that the embodied carbon in imports is smaller, but still very significant: about 6.2 tons CO₂e per year per person. In energy terms, 6 tons CO₂e per year is something like 60 kWh/d.

Here, let's see if we can reproduce these conclusions in a different way, using the weights of the imports.

Figure H.2 shows Britain's imports in the year 2006 in three ways: on the left, the total *value* of the imports is broken down by the country of origin. In the middle, the same total financial value is broken down by the type of stuff imported, using the categories of HM Revenue and Customs. On the right, all maritime imports to Britain are shown by *weight* and broken down by the categories used by the Department for Transport, which doesn't care whether something is leather or tobacco – it keeps track of how heavy stuff is, whether it is dry or liquid, and whether the stuff arrived in a container or a lorry.

The energy cost of the imported fuels (top right) *is* included in the standard accounts of British energy consumption; the energy costs of all the other imports are not. For most materials, the embodied energy per unit weight is greater than or equal to 10 kWh per kg – the same as the energy per unit weight of fossil fuels. This is true of all metals and alloys, all polymers and composites, most paper products, and many ceramics, for example. The exceptions are raw materials like ores; porous ceramics such as concrete, brick, and porcelain, whose energy cost is 10 times lower; wood and some rubbers; and glasses, whose energy cost is a whisker lower than 10 kWh per kg. [r22oz]

We can thus roughly estimate the energy footprint of our imports simply from the weight of their manufactured materials, if we exclude things like ores and wood. Given the crudity of the data with which we are working, we will surely slip up and inadvertently include some things made of wood and glass, but hopefully such slips will be balanced by our underestimation of the energy content of most of the metals and plastics and more complex goods, many of which have an embodied energy of not 10 but 30 kWh per kg, or even more.

For this calculation I'll take from the right-hand column in figure H.2




Figure H.1. Continuous casting of steel strands at Korea Iron and Steel Company.




Figure H.2. Imports of stuff to the UK, 2006.

the iron and steel products, the dry bulk products, the containerized freight and the “other freight,” which total 98 million tons per year. I’m leaving the vehicles to one side for a moment. I subtract from this an estimated 25 million tons of food which is presumably lurking in the “other freight” category (34 million tons of food were imported in 2006), leaving 73 million tons.

Converting 73 million tons to energy using the exchange rate suggested above, and sharing between 60 million people, we estimate that those imports have an embodied energy of 33 kWh/d per person.

For the cars, we can hand-wave a little less, because we know a little more: the number of imported vehicles in 2006 was 2.4 million. If we take the embodied energy per car to be 76 000 kWh (a number we picked up on p90) then these imported cars have an embodied energy of 8 kWh/d per person.

I left the “liquid bulk products” out of these estimates because I am not sure what sort of products they are. If they are actually liquid chemicals then their contribution might be significant.

We’ve arrived at a total estimate of 41 kWh/d per person for the embodied energy of imports – definitely in the same ballpark as the estimate of Dieter Helm and his colleagues.

I suspect that 41 kWh/d per person may be an underestimate because the energy intensity we assumed (10 kWh/d per person) is too low for most forms of manufactured goods such as machinery or electrical equipment. However, without knowing the weights of all the import categories, this is the best estimate I can make for now.



Figure H.3. Niobium open cast mine, Brazil.

Lifecycle analysis for buildings

Tables H.4 and H.5 show estimates of the *Process Energy Requirement* of building materials and building constructions. This includes the energy used in transporting the raw materials to the factory but not energy used to transport the final product to the building site.

Table H.6 uses these numbers to estimate the process energy for making a three-bedroom house. The *gross energy requirement* widens the boundary, including the embodied energy of urban infrastructure, for example, the embodied energy of the machinery that makes the raw materials. A rough rule of thumb to get the gross energy requirement of a building is to double the process energy requirement [3kmcks].

If we share 42 000 kWh over 100 years, and double it to estimate the gross energy cost, the total embodied energy of a house comes to about 2.3 kWh/d. This is the energy cost of the *shell* of the house only – the bricks, tiles, roof beams.

Material	Embodied energy	
	(MJ/kg)	(kWh/kg)
kiln-dried sawn softwood	3.4	0.94
kiln-dried sawn hardwood	2.0	0.56
air dried sawn hardwood	0.5	0.14
hardboard	24.2	6.7
particleboard	8.0	2.2
MDF	11.3	3.1
plywood	10.4	2.9
glue-laminated timber	11	3.0
laminated veneer lumber	11	3.0
straw	0.24	0.07
stabilised earth	0.7	0.19
imported dimension granite	13.9	3.9
local dimension granite	5.9	1.6
gypsum plaster	2.9	0.8
plasterboard	4.4	1.2
fibre cement	4.8	1.3
cement	5.6	1.6
in situ concrete	1.9	0.53
precast steam-cured concrete	2.0	0.56
precast tilt-up concrete	1.9	0.53
clay bricks	2.5	0.69
concrete blocks	1.5	0.42
autoclaved aerated concrete	3.6	1.0
plastics – general	90	25
PVC	80	22
synthetic rubber	110	30
acrylic paint	61.5	17
glass	12.7	3.5
fibreglass (glasswool)	28	7.8
aluminium	170	47
copper	100	28
galvanised steel	38	10.6
stainless steel	51.5	14.3

Table H.4. Embodied energy of building materials (assuming virgin rather than recycled product is used). (Dimension stone is natural stone or rock that has been selected and trimmed to specific sizes or shapes.) Sources: [3kmcks], Lawson (1996).



	Embodied energy (kWh/m ²)
Walls	
timber frame, timber weatherboard, plasterboard lining	52
timber frame, clay brick veneer, plasterboard lining	156
timber frame, aluminium weatherboard, plasterboard lining	112
steel frame, clay brick veneer, plasterboard lining	168
double clay brick, plasterboard lined	252
cement stabilised rammed earth	104
Floors	
elevated timber floor	81
110 mm concrete slab on ground	179
200 mm precast concrete T beam/infill	179
Roofs	
timber frame, concrete tile, plasterboard ceiling	70
timber frame, terracotta tile, plasterboard ceiling	75
timber frame, steel sheet, plasterboard ceiling	92

Table H.5. Embodied energy in various walls, floors, and roofs. Sources: [3kmcks], Lawson (1996).

	Area (m ²)	×	energy density (kWh/m ²)	=	energy (kWh)
Floors	100	×	81	=	8100
Roof	75	×	75	=	5600
External walls	75	×	252	=	19 000
Internal walls	75	×	125	=	9400
Total					42 000

Table H.6. Process energy for making a three-bedroom house.

Notes and further reading

page no.

322 A subsequent more-detailed study commissioned by DEFRA estimated that the embodied carbon in imports is about 6.2 tons CO₂e per person. Wiedmann et al. (2008).

Further resources: www.greenbooklive.com has life cycle assessments of building products.
Some helpful cautions about life-cycle analysis: www.gdrc.org/uem/lca/life-cycle.html.
More links: www.epa.gov/ord/NRMRL/lcaccess/resources.htm.



Figure H.7. Millau Viaduct in France, the highest bridge in the world. Steel and concrete, 2.5 km long and 353 m high.